Cisco IE 2000 Switch Software Configuration Guide

Cisco IOS Release 15.0(1)EY
July 2012

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Text Part Number: OL-25866-01

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco IE 2000 Switch Software Configuration Guide
© 2012 Cisco Systems, Inc. All rights reserved.

CONTENTS

Preface li

Audience li

Purpose li

Conventions li

Related Publications liii

Obtaining Documentation, Obtaining Support, and Security Guidelines liii

CHAPTER 1 **Configuration Overview** 1-1

Features 1-1

Feature Software Licensing 1-1

Ease-of-Deployment and Ease-of-Use Features 1-2

Performance Features 1-2

Management Options 1-3

Industrial Application 1-4

Manageability Features 1-4

Availability and Redundancy Features 1-5

VLAN Features 1-6

Security Features 1-7

QoS and CoS Features 1-10

Monitoring Features 1-11

Default Settings After Initial Switch Configuration 1-11

Network Configuration Examples 1-14

Design Concepts for Using the Switch 1-14

Ethernet-to-the-Factory Architecture 1-15

Enterprise Zone 1-15

Demilitarized Zone 1-16

Manufacturing Zone 1-16

Topology Options 1-18

Where to Go Next 1-21

CHAPTER 2 **Using the Command-Line Interface** 2-1

Information About Using the Command-Line Interface 2-1

Command Modes 2-1

Help System 2-3

Understanding Abbreviated Commands 2-4
No and default Forms of Commands 2-4
CLI Error Messages 2-5
Configuration Logging 2-5
How to Use the CLI to Configure Features 2-6
Configuring the Command History 2-6
Changing the Command History Buffer Size 2-6
Recalling Commands 2-6
Disabling the Command History Feature 2-7
Using Editing Features 2-7
Enabling and Disabling Editing Features 2-7
Editing Commands Through Keystrokes 2-7
Editing Command Lines That Wrap 2-9
Searching and Filtering Output of show and more Commands 2-10
Accessing the CLI 2-10
Accessing the CLI through a Console Connection or through Telnet 2-10

CHAPTER 3 **Configuring Switch Alarms** 3-1
Finding Feature Information 3-1
Information About Switch Alarms 3-1
Global Status Monitoring Alarms 3-2
FCS Error Hysteresis Threshold 3-2
Port Status Monitoring Alarms 3-2
Triggering Alarm Options 3-3
External Alarms 3-4
Default Switch Alarm Settings 3-5
How to Configure Switch Alarms 3-5
Configuring External Alarms 3-5
Configuring the Power Supply Alarms 3-6
Configuring the Switch Temperature Alarms 3-6
Associating the Temperature Alarms to a Relay 3-7
Configuring the FCS Bit Error Rate Alarm 3-7
Setting the FCS Error Threshold 3-7
Setting the FCS Error Hysteresis Threshold 3-8
Configuring Alarm Profiles 3-8
Creating an Alarm Profile 3-8
Modifying an Alarm Profile 3-8
Attaching an Alarm Profile to a Specific Port 3-9
Enabling SNMP Traps 3-9

Monitoring and Maintaining Switch Alarms Status 3-9
Configuration Examples for Switch Alarms 3-10
Configuring External Alarms: Example 3-10
Associating Temperature Alarms to a Relay: Examples 3-10
Creating or Modifying an Alarm Profile: Example 3-10
Setting the FCS Error Hysteresis Threshold: Example 3-11
Configuring a Dual Power Supply: Examples 3-11
Displaying Alarm Settings: Example 3-11
Additional References 3-12
Related Documents 3-12
Standards 3-12
MIBs 3-12
RFCs 3-13
Technical Assistance 3-13

CHAPTER 4 **Performing Switch Setup Configuration** 4-1
Restrictions for Performing Switch Setup Configuration 4-1
Information About Performing Switch Setup Configuration 4-1
Switch Boot Process 4-1
Default Switch Boot Settings 4-3
Switch Boot Optimization 4-3
Switch Information Assignment 4-4
Switch Default Settings 4-4
DHCP-Based Autoconfiguration Overview 4-4
DHCP Client Request Process 4-5
DHCP-Based Autoconfiguration and Image Update 4-6
DHCP Autoconfiguration 4-6
DHCP Auto-Image Update 4-6
DHCP Server Configuration Guidelines 4-7
TFTP Server 4-7
DNS Server 4-8
Relay Device 4-8
How to Obtain Configuration Files 4-9
How to Control Environment Variables 4-10
Common Environment Variables 4-11
Scheduled Reload of the Software Image 4-11
How to Perform Switch Setup Configuration 4-12
Configuring DHCP Autoconfiguration (Only Configuration File) 4-12
Configuring DHCP Auto-Image Update (Configuration File and Image) 4-13

Configuring the Client 4-14
Manually Assigning IP Information on a Routed Port 4-14
Manually Assigning IP Information to SVIs 4-15
Modifying the Startup Configuration 4-15
Specifying the Filename to Read and Write the System Configuration 4-15
Manually Booting the Switch 4-16
Booting a Specific Software Image 4-17
Monitoring Switch Setup Configuration 4-17
Verifying the Switch Running Configuration 4-17
Configuration Examples for Performing Switch Setup Configuration 4-18
Retrieving IP Information Using DHCP-Based Autoconfiguration: Example 4-18
Scheduling Software Image Reload: Examples 4-20
Configuring DHCP Auto-Image Update: Example 4-20
Configuring a Switch as a DHCP Server: Example 4-20
Configuring Client to Download Files from DHCP Server 4-21
Additional References 4-22
Related Documents 4-22
Standards 4-22
MIBs 4-22
RFCs 4-22
Technical Assistance 4-22

CHAPTER 5 **Configuring Cisco IOS Configuration Engine 5-1**
Finding Feature Information 5-1
Prerequisites for Configuring Cisco IOS Configuration Engine 5-1
Information About Configuring Cisco IOS Configuration Engine 5-2
Configuration Service 5-3
Event Service 5-3
NameSpace Mapper 5-4
CNS IDs and Device Hostnames 5-4
ConfigID 5-4
DeviceID 5-4
Hostname and DeviceID Interaction 5-5
Using Hostname, DeviceID, and ConfigID 5-5
Cisco IOS Agents 5-5
Initial Configuration 5-5
Incremental (Partial) Configuration 5-6
Synchronized Configuration 5-6
How to Configure Cisco IOS Configuration Engine 5-7

Configuring Cisco IOS Agents 5-7
Enabling CNS Event Agent 5-7
Enabling Cisco IOS CNS Agent and an Initial Configuration 5-8
Enabling a Partial Configuration 5-10
Monitoring and Maintaining Cisco IOS Configuration Engine 5-11
Configuration Examples for Cisco IOS Configuration Engine 5-11
Enabling the CNS Event Agent: Example 5-11
Configuring an Initial CNS Configuration: Examples 5-11
Additional References 5-12
Related Documents 5-12
Standards 5-12
MIBs 5-12
RFCs 5-12
Technical Assistance 5-13

CHAPTER 6 **Configuring Switch Clusters** 6-1
Finding Feature Information 6-1
Prerequisites for Configuring Switch Clusters 6-1
Cluster Command Switch Characteristics 6-1
Standby Cluster Command Switch Characteristics 6-2
Candidate Switch and Cluster Member Switch Characteristics 6-2
Restrictions for Configuring Switch Clusters 6-3
Information About Configuring Switch Clusters 6-3
Benefits of Clustering Switches 6-3
Eligible Cluster Switches 6-3
How to Plan for Switch Clustering 6-4
Automatic Discovery of Cluster Candidates and Members 6-5
Discovery Through CDP Hops 6-5
Discovery Through Non-CDP-Capable and Noncluster-Capable Devices 6-7
Discovery Through Different VLANs 6-7
Discovery Through Different Management VLANs 6-8
Discovery Through Routed Ports 6-9
Discovery of Newly Installed Switches 6-10
IP Addresses 6-11
Hostnames 6-11
Passwords 6-12
SNMP Community Strings 6-12
TACACS+ and RADIUS 6-12
LRE Profiles 6-13

Managing Switch Clusters 6-13
Using the CLI to Manage Switch Clusters 6-13
Using SNMP to Manage Switch Clusters 6-14
Additional References 6-15
Related Documents 6-15
Standards 6-15
MIBs 6-15
RFCs 6-15
Technical Assistance 6-15

CHAPTER 7 **Performing Switch Administration** 7-1
Finding Feature Information 7-1
Information About Performing Switch Administration 7-1
System Time and Date Management 7-1
System Clock 7-1
Network Time Protocol 7-2
NTP Version 4 7-3
DNS 7-4
Default DNS Configuration 7-4
Login Banners 7-4
System Name and Prompt 7-5
MAC Address Table 7-5
Address Table 7-5
MAC Addresses and VLANs 7-5
Default MAC Address Table Configuration 7-6
Address Aging Time for VLANs 7-6
MAC Address Change Notification Traps 7-6
Static Addresses 7-6
Unicast MAC Address Filtering 7-7
MAC Address Learning on a VLAN 7-8
ARP Table Management 7-8
How to Perform Switch Administration 7-9
Configuring Time and Date Manually 7-9
Setting the System Clock 7-9
Configuring the Time Zone 7-9
Configuring Summer Time (Daylight Saving Time) 7-10
Configuring Summer Time (Exact Date and Time) 7-11
Configuring a System Name 7-11
Setting Up DNS 7-11

Configuring Login Banners 7-12
Configuring a Message-of-the-Day Login Banner 7-12
Configuring a Login Banner 7-13
Managing the MAC Address Table 7-13
Changing the Address Aging Time 7-13
Configuring MAC Address Change Notification Traps 7-14
Configuring MAC Address Move Notification Traps 7-15
Configuring MAC Threshold Notification Traps 7-15
Adding and Removing Static Address Entries 7-17
Configuring Unicast MAC Address Filtering 7-17
Disabling MAC Address Learning on a VLAN 7-17
Monitoring and Maintaining Switch Administration 7-18
Configuration Examples for Performing Switch Admininistration 7-18
Setting the System Clock: Example 7-18
Configuring Summer Time: Examples 7-18
Configuring a MOTD Banner: Examples 7-19
Configuring a Login Banner: Example 7-19
Configuring MAC Address Change Notification Traps: Example 7-19
Sending MAC Address Move Notification Traps: Example 7-20
Configuring MAC Threshold Notification Traps: Example 7-20
Adding the Static Address to the MAC Address Table: Example 7-20
Configuring Unicast MAC Address Filtering: Example 7-20
Additional References 7-21
Related Documents 7-21
Standards 7-21
MIBs 7-21
RFCs 7-21
Technical Assistance 7-21

CHAPTER 8 **Configuring PTP** 8-1
Finding Feature Information 8-1
Prerequisites for Configuring PTP 8-1
Restrictions for Configuring PTP 8-1
Information About Configuring PTP 8-1
Precision Time Protocol 8-1
How to Configure PTP 8-2
Default PTP Settings 8-2
Setting Up PTP 8-3
Monitoring and Maintaining the PTP Configuration 8-3

Troubleshooting the PTP Configuration 8-4
Additional References 8-4
Related Documents 8-4
Standards 8-4
MIBs 8-4
RFCs 8-5
Technical Assistance 8-5

CHAPTER 9 **Configuring PROFINET** 9-1
Finding Feature Information 9-1
Restrictions for Configuring PROFINET 9-1
Information About Configuring PROFINET 9-1
PROFINET Device Roles 9-2
PROFINET Device Data Exchange 9-2
How to Configure PROFINET 9-4
Configuring PROFINET 9-4
Default Configuration 9-4
Enabling PROFINET 9-4
Monitoring and Maintaining PROFINET 9-5
Troubleshooting PROFINET 9-5
Additional References 9-6
Related Documents 9-6
Standards 9-6
MIBs 9-6
RFCs 9-6
Technical Assistance 9-6

CHAPTER 10 **Configuring CIP** 10-1
Finding Feature Information 10-1
Restrictions for Configuring CIP 10-1
Information About Configuring CIP 10-1
How to Configure CIP 10-1
Default Configuration 10-1
Enabling CIP 10-2
Monitoring CIP 10-2
Troubleshooting CIP 10-2
Additional References 10-3
Related Documents 10-3

Standards 10-3
MIBs 10-3
RFCs 10-3
Technical Assistance 10-3

CHAPTER 11 **Configuring SDM Templates** 11-1
Finding Feature Information 11-1
Prerequisites for Configuring SDM Templates 11-1
Restrictions for Configuring SDM Templates 11-1
Information About Configuring SDM Templates 11-1
SDM Templates 11-1
Dual IPv4 and IPv6 SDM Default Template 11-3
How to Configure the Switch SDM Templates 11-4
Setting the SDM Template 11-4
Monitoring and Maintaining SDM Templates 11-4
Configuration Examples for Configuring SDM Templates 11-5
Configuring the IPv4-and-IPv6 Default Template: Example 11-5
Additional References 11-6
Related Documents 11-6
Standards 11-6
MIBs 11-6
RFCs 11-6
Technical Assistance 11-6

CHAPTER 12 **Configuring Switch-Based Authentication** 12-1
Finding Feature Information 12-1
Prerequisites for Configuring Switch-Based Authentication 12-1
Restrictions for Configuring Switch-Based Authentication 12-1
Information About Configuring Switch-Based Authentication 12-2
Prevention for Unauthorized Switch Access 12-2
Password Protection 12-2
Default Password and Privilege Level Configuration 12-2
Enable Secret Passwords with Encryption 12-3
Password Recovery 12-3
Telnet Password for a Terminal Line 12-4
Username and Password Pairs 12-4
Multiple Privilege Levels 12-4
Switch Access with TACACS+ 12-5

TACACS+ 12-5
TACACS+ Operation 12-6
Default TACACS+ Configuration 12-7
TACACS+ Server Host and the Authentication Key 12-7
TACACS+ Login Authentication 12-7
TACACS+ Authorization for Privileged EXEC Access and Network Services 12-7
TACACS+ Accounting 12-8
Switch Access with RADIUS 12-8
RADIUS 12-8
RADIUS Operation 12-9
Default RADIUS Configuration 12-10
RADIUS Change of Authorization 12-10
CoA Request Commands 12-12
RADIUS Server Host 12-14
RADIUS Login Authentication 12-15
Radius Method List 12-15
AAA Server Groups 12-15
RADIUS Authorization for User Privileged Access and Network Services 12-16
RADIUS Accounting 12-16
Establishing a Session with a Router if the AAA Server is Unreachable 12-16
Vendor-Specific RADIUS Attributes 12-16
Vendor-Proprietary RADIUS Server Communication 12-17
Switch Access with Kerberos 12-17
Understanding Kerberos 12-17
Kerberos Operation 12-19
Kerberos Configuration 12-20
Local Authentication and Authorization 12-20
Secure Shell 12-21
SSH 12-21
SSH Servers, Integrated Clients, and Supported Versions 12-21
Limitations 12-22
SSH Configuration Guidelines 12-22
Switch for Secure Socket Layer HTTP 12-22
Secure HTTP Servers and Clients 12-22
Default SSL Settings 12-23
Certificate Authority Trustpoints 12-23
CipherSuites 12-24
Secure Copy Protocol 12-24
How to Configure Switch-Based Authentication 12-26
Configuring Password Protection 12-26

Setting or Changing a Static Enable Password 12-26
Protecting Enable and Enable Secret Passwords with Encryption 12-27
Disabling Password Recovery 12-27
Setting a Telnet Password for a Terminal Line 12-28
Configuring Username and Password Pairs 12-28
Setting the Privilege Level for a Command 12-29
Changing the Default Privilege Level for Lines 12-29
Logging Into and Exiting a Privilege Level 12-30
Configuring TACACS+ 12-30
Identifying the TACACS+ Server Host and Setting the Authentication Key 12-30
Configuring TACACS+ Login Authentication 12-31
Configuring TACACS+ Authorization for Privileged EXEC Access and Network Services 12-33
Starting TACACS+ Accounting 12-33
Configuring Radius Server Communication 12-33
Defining AAA Server Groups 12-35
Configuring RADIUS Login Authentication 12-36
Configuring RADIUS Authorization for User Privileged Access and Network Services 12-37
Starting RADIUS Accounting 12-37
Configuring Settings for All RADIUS Servers 12-37
Configuring the Switch for Vendor-Proprietary RADIUS Server Communication 12-38
Configuring CoA on the Switch 12-38
Configuring the Switch for Local Authentication and Authorization 12-39
Configuring Secure Shell 12-40
Setting Up the Switch to Run SSH 12-40
Configuring the SSH Server 12-40
Configuring Secure HTTP Servers and Clients 12-42
Configuring a CA Trustpoint 12-42
Configuring the Secure HTTP Server 12-42
Configuring the Secure HTTP Client 12-44
Monitoring and Maintaining Switch-Based Authentication 12-44
Configuration Examples for Configuring Switch-Based Authentication 12-45
Changing the Enable Password: Example 12-45
Configuring the Encrypted Password: Example 12-45
Setting the Telnet Password for a Terminal Line: Example 12-45
Setting the Privilege Level for a Command: Example 12-45
Configuring the RADIUS Server: Examples 12-45
Defining AAA Server Groups: Example 12-46
Configuring Vendor-Specific RADIUS Attributes: Examples 12-46
Configuring a Vendor-Proprietary RADIUS Host: Example 12-46
Sample Output for a Self-Signed Certificate: Example 12-46

Verifying Secure HTTP Connection: Example 12-47
Additional References 12-47
Related Documents 12-47
Standards 12-48
MIBs 12-48
RFCs 12-48
Technical Assistance 12-48

CHAPTER 13 **Configuring IEEE 802.1x Port-Based Authentication** 13-1
Finding Feature Information 13-1
Restrictions for Configuring IEEE 802.1x Port-Based Authentication 13-1
Information About Configuring IEEE 802.1x Port-Based Authentication 13-1
IEEE 802.1x Port-Based Authentication 13-1
Device Roles 13-2
Authentication Process 13-3
Switch-to-RADIUS-Server Communication 13-4
Authentication Initiation and Message Exchange 13-4
Authentication Manager 13-6
Port-Based Authentication Methods 13-6
Per-User ACLs and Filter-Ids 13-7
Authentication Manager CLI Commands 13-8
Ports in Authorized and Unauthorized States 13-9
802.1x Host Mode 13-9
Multidomain Authentication 13-10
802.1x Multiple Authentication Mode 13-11
MAC Move 13-12
MAC Replace 13-12
802.1x Accounting 13-13
802.1x Accounting Attribute-Value Pairs 13-13
802.1x Readiness Check 13-14
802.1x Authentication with VLAN Assignment 13-15
Voice Aware 802.1x Security 13-16
802.1x Authentication with Per-User ACLs 13-17
802.1x Authentication with Downloadable ACLs and Redirect URLs 13-18
Cisco Secure ACS and Attribute-Value Pairs for the Redirect URL 13-19
Cisco Secure ACS and Attribute-Value Pairs for Downloadable ACLs 13-19
VLAN ID-Based MAC Authentication 13-20
802.1x Authentication with Guest VLAN 13-20
802.1x Authentication with Restricted VLAN 13-21

802.1x Authentication with Inaccessible Authentication Bypass 13-22
Support on Multiple-Authentication Ports 13-22
Authentication Results 13-22
Feature Interactions 13-23
802.1x Authentication with Voice VLAN Ports 13-23
802.1x Authentication with Port Security 13-24
802.1x Authentication with Wake-on-LAN 13-24
802.1x Authentication with MAC Authentication Bypass 13-25
802.1x User Distribution 13-26
802.1x User Distribution Configuration Guidelines 13-26
Network Admission Control Layer 2 802.1x Validation 13-27
Flexible Authentication Ordering 13-27
Open1x Authentication 13-28
802.1x Supplicant and Authenticator Switches with Network Edge Access Topology (NEAT) 13-28
802.1x Supplicant and Authenticator Switch Guidelines 13-29
Using IEEE 802.1x Authentication with ACLs and the RADIUS Filter-Id Attribute 13-29
Authentication Manager Common Session ID 13-30
Default 802.1x Authentication Settings 13-30
802.1x Accounting 13-31
802.1x Authentication Guidelines 13-32
VLAN Assignment, Guest VLAN, Restricted VLAN, and Inaccessible Authentication Bypass Guidelines 13-33
MAC Authentication Bypass Guidelines 13-33
Maximum Number of Allowed Devices Per Port Guidelines 13-34
How to Configure IEEE 802.1x Port-Based Authentication 13-34
802.1x Authentication Configuration Process 13-34
Configuring the Switch-to-RADIUS-Server Communication 13-36
Configuring 802.1x Readiness Check 13-36
Enabling Voice Aware 802.1x Security 13-37
Configuring 802.1x Violation Modes 13-37
Configuring the Host Mode 13-38
Configuring Periodic Reauthentication 13-39
Configuring Optional 802.1x Authentication Features 13-40
Configuring 802.1x Accounting 13-42
Configuring a Guest VLAN 13-42
Configuring a Restricted VLAN 13-43
Configuring the Maximum Number of Authentication Attempts 13-43
Configuring Inaccessible Authentication Bypass 13-44
Configuring 802.1x User Distribution 13-46
Configuring NAC Layer 2 802.1x Validation 13-46

Configuring an Authenticator and Supplicant 13-47
Configuring an Authenticator 13-47
Configuring a Supplicant Switch with NEAT 13-47
Configuring 802.1x Authentication with Downloadable ACLs and Redirect URLs 13-48
Configuring Downloadable ACLs 13-48
Configuring a Downloadable Policy 13-49
Configuring Open1x 13-50
Resetting the 802.1x Authentication Configuration to the Default Values 13-51
Monitoring and Maintaining IEEE 802.1x Port-Based Authentication 13-51
Configuration Examples for Configuring IEEE 802.1x Port-Based Authentication 13-51
Enabling a Readiness Check: Example 13-51
Enabling 802.1x Authentication: Example 13-52
Enabling MDA: Example 13-52
Disabling the VLAN Upon Switch Violoation: Example 13-52
Configuring the Radius Server Parameters: Example 13-52
Configuring 802.1x Accounting: Example 13-52
Enabling an 802.1x Guest VLAN: Example 13-53
Displaying Authentication Manager Common Session ID: Examples 13-53
Configuring Inaccessible Authentication Bypass: Example 13-53
Configuring VLAN Groups: Examples 13-54
Configuring NAC Layer 2 802.1x Validation: Example 13-54
Configuring an 802.1x Authenticator Switch: Example 13-54
Configuring an 802.1x Supplicant Switch: Example 13-55
Configuring a Downloadable Policy: Example 13-55
Configuring Open 1x on a Port: Example 13-55
Additional References 13-56
Related Documents 13-56
Standards 13-56
MIBs 13-56
RFCs 13-56
Technical Assistance 13-57

CHAPTER 14 **Configuring Web-Based Authentication** 14-1
Finding Feature Information 14-1
Prerequisites for Configuring Web-Based Authentication 14-1
Restrictions for Configuring Web-Based Authentication on the IE 2000 Switch 14-1
Information About Configuring Web-Based Authentication 14-2
Web-Based Authentication 14-2
Device Roles 14-2

Host Detection 14-3
Session Creation 14-3
Authentication Process 14-4
Local Web Authentication Banner 14-4
Web Authentication Customizable Web Pages 14-6
Web Authentication Guidelines 14-6
Web-Based Authentication Interactions with Other Features 14-8
Port Security 14-8
LAN Port IP 14-8
Gateway IP 14-9
ACLs 14-9
Context-Based Access Control 14-9
802.1x Authentication 14-9
EtherChannel 14-9
Default Web-Based Authentication Settings 14-10
Configuring Switch-to-RADIUS-Server Communication 14-10
How to Configure Web-Based Authentication 14-11
Configuring the Authentication Rule and Interfaces 14-11
Configuring AAA Authentication 14-11
Configuring Switch-to-RADIUS-Server Communication 14-12
Configuring the HTTP Server 14-12
Customizing the Authentication Proxy Web Pages 14-13
Specifying a Redirection URL for Successful Login 14-13
Configuring the Web-Based Authentication Parameters 14-13
Configuring a Web Authentication Local Banner 14-14
Removing Web-Based Authentication Cache Entries 14-14
Monitoring and Maintaining Web-Based Authentication 14-14
Configuration Examples for Configuring Web-Based Authentication 14-14
Enabling and Displaying Web-Based Authentication: Examples 14-14
Enabling AAA: Example 14-15
Configuring the RADIUS Server Parameters: Example 14-15
Configuring a Custom Authentication Proxy Web Page: Example 14-15
Verifying a Custom Authentication Proxy Web Page: Example 14-15
Configuring a Redirection URL: Example 14-16
Verifying a Redirection URL: Example 14-16
Configuring a Local Banner: Example 14-16
Clearing the Web-Based Authentication Session: Example 14-16
Additional References 14-17
Related Documents 14-17

Standards 14-17
MIBs 14-17
RFCs 14-18
Technical Assistance 14-18

CHAPTER 15 **Configuring Interface Characteristics 15-1**
Finding Feature Information 15-1
Restrictions for Configuring Interface Characteristics 15-1
Information About Configuring Interface Characteristics 15-1
Interface Types 15-1
Port-Based VLANs 15-2
Switch Ports 15-2
Routed Ports 15-3
Access Ports 15-3
Trunk Ports 15-4
EtherChannel Port Groups 15-4
Dual-Purpose Uplink Ports 15-4
Connecting Interfaces 15-5
Using Interface Configuration Mode 15-6
Default Ethernet Interface Settings 15-8
Interface Speed and Duplex Mode 15-9
Speed and Duplex Configuration Guidelines 15-9
IEEE 802.3x Flow Control 15-9
Auto-MDIX on an Interface 15-10
SVI Autostate Exclude 15-10
System MTU 15-10
How to Configure Interface Characteristics 15-11
Configuring Layer 3 Interfaces 15-11
Configuring Interfaces 15-13
Configuring a Range of Interfaces 15-13
Interface Range Restrictions 15-13
Configuring and Using Interface Range Macros 15-14
Configuring Ethernet Interfaces 15-15
Setting the Type of a Dual-Purpose Uplink Port 15-15
Setting the Interface Speed and Duplex Parameters 15-16
Configuring IEEE 802.3x Flow Control 15-16
Configuring Auto-MDIX on an Interface 15-17
Adding a Description for an Interface 15-17
Configuring SVI Autostate Exclude 15-17

Configuring the System MTU 15-18
Monitoring and Maintaining Interface Characteristics 15-18
Monitoring Interface Status 15-18
Clearing and Resetting Interfaces and Counters 15-19
Shutting Down and Restarting the Interface 15-19
Configuration Examples for Configuring Interface Characteristics 15-20
Configuring the Interface Range: Examples 15-20
Configuring Interface Range Macros: Examples 15-20
Setting Speed and Duplex Parameters: Example 15-21
Enabling auto-MDIX: Example 15-21
Adding a Description on a Port: Example 15-21
Configuring SVI Autostate Exclude: Example 15-22
Additional References 15-22
Related Documents 15-22
Standards 15-22
MIBs 15-22
RFCs 15-23

CHAPTER 16 **Configuring Smartports Macros** 16-1
Finding Feature Information 16-1
Information About Configuring Smartports Macros 16-1
How to Configure Smartports Macros 16-1
Default Smartports Settings 16-1
Smartports Configuration Guidelines 16-2
Applying Smartports Macros 16-3
Monitoring and Maintaining Smartports Macros 16-4
Configuration Examples for Smartports Macros 16-4
Applying the Smartports Macro: Examples 16-4
Additional References 16-5
Related Documents 16-5
Standards 16-5
MIBs 16-5
RFCs 16-6
Technical Assistance 16-6

CHAPTER 17 **Configuring VLANs** 17-1
Finding Feature Information 17-1
Information About Configuring VLANs 17-1

VLANs 17-1
Supported VLANs 17-2
VLAN Port Membership Modes 17-3
Normal-Range VLANs 17-4
Token Ring VLANs 17-5
Normal-Range VLAN Configuration Guidelines 17-6
Default Ethernet VLAN Configuration 17-6
Ethernet VLANs 17-7
VLAN Removal 17-7
Static-Access Ports for a VLAN 17-7
Extended-Range VLANs 17-8
Default VLAN Configuration 17-8
Extended-Range VLAN Configuration Guidelines 17-8
VLAN Trunks 17-9
Trunking Overview 17-9
IEEE 802.1Q Configuration Guidelines 17-10
Default Layer 2 Ethernet Interface VLAN Settings 17-11
Ethernet Interface as a Trunk Port 17-11
Trunking Interaction with Other Features 17-11
Allowed VLANs on a Trunk 17-12
Native VLAN for Untagged Traffic 17-12
Load Sharing Using Trunk Ports 17-12
Load Sharing Using STP Port Priorities 17-13
Load Sharing Using STP Path Cost 17-13
VMPS 17-14
Dynamic-Access Port VLAN Membership 17-15
Default VMPS Client Settings 17-15
VMPS Configuration Guidelines 17-15
VMPS Reconfirmation Interval 17-16
Dynamic-Access Port VLAN Membership 17-16
How to Configure VLANs 17-17
Creating or Modifying an Ethernet VLAN 17-17
Deleting a VLAN 17-17
Assigning Static-Access Ports to a VLAN 17-17
Creating an Extended-Range VLAN 17-18
Creating an Extended-Range VLAN with an Internal VLAN ID 17-18
Configuring an Ethernet Interface as a Trunk Port 17-19
Defining the Allowed VLANs on a Trunk 17-19
Changing the Pruning-Eligible List 17-19
Configuring the Native VLAN for Untagged Traffic 17-20

Load Sharing Using STP Port Priorities 17-21
Configuring Load Sharing Using STP Path Cost 17-21
Configuring the VMPS Client 17-22
Entering the IP Address of the VMPS 17-22
Configuring Dynamic-Access Ports on VMPS Clients 17-23
Monitoring and Maintaining VLANs 17-23
Configuration Examples for Configuring VLANs 17-24
VMPS Network: Example 17-24
Configuring a VLAN: Example 17-25
Configuring an Access Port in a VLAN: Example 17-25
Configuring an Extended-Range VLAN: Example 17-25
Configuring a Trunk Port: Example 17-25
Removing a VLAN: Example 17-25
Show VMPS Output: Example 17-25
Additional References 17-26
Related Documents 17-26
Standards 17-26
MIBs 17-26
RFCs 17-26

CHAPTER 18 **Configuring VTP** 18-1
Finding VTP Feature Information 18-1
Prerequisites for Configuring VTP 18-1
Restrictions for Configuring VTP 18-1
Information About Configuring VTP 18-2
VTP 18-2
VTP Domain 18-2
VTP Modes 18-3
VTP Mode Guidelines 18-3
VTP Advertisements 18-4
VTP Version 2 18-5
VTP Version 3 18-5
VTP Version Guidelines 18-6
VTP Pruning 18-7
Default VTP Settings 18-9
VTP Configuration Guidelines 18-9
Domain Names 18-10
Passwords 18-10
Adding a VTP Client Switch to a VTP Domain 18-10

How to Configure VTP 18-11
Configuring VTP Domain and Parameters 18-11
Configuring a VTP Version 3 Password 18-12
Enabling the VTP Version 18-12
Enabling VTP Pruning 18-13
Configuring VTP on a Per-Port Basis 18-13
Adding a VTP Client Switch to a VTP Domain 18-13
Monitoring and Maintaining VTP 18-14
Configuration Examples for Configuring VTP 18-14
Configuring a VTP Server: Example 18-14
Configuring a Hidden VTP Password: Example 18-15
Configuring a VTP Version 3 Primary Server: Example 18-15
Additional References for Configuring VTP 18-15
Related Documents 18-15
Standards 18-15
MIBs 18-16
RFCs 18-16

CHAPTER 19 **Configuring Voice VLAN** 19-1
Finding Feature Information 19-1
Information About Configuring Voice VLAN 19-1
Voice VLAN 19-1
Cisco IP Phone Voice Traffic 19-2
Cisco IP Phone Data Traffic 19-3
Default Voice VLAN Configuration 19-3
Voice VLAN Configuration Guidelines 19-3
Port Connection to a Cisco 7960 IP Phone 19-4
Priority of Incoming Data Frames 19-4
How to Configure VTP 19-5
Configuring Cisco IP Phone for Voice Traffic 19-5
Configuring the Priority of Incoming Data Frames 19-5
Monitoring and Maintaining Voice VLAN 19-6
Configuration Examples for Configuring Voice VLAN 19-6
Configuring a Cisco IP Phone for Voice Traffic: Example 19-6
Configuring the Cisco IP Phone Priority of Incoming Data Frames: Example 19-6
Additional References for Configuring Voice VLAN 19-6
Related Documents 19-6
Standards 19-7
MIBs 19-7

RFCs 19-7

CHAPTER 20 **Configuring STP** 20-1

Finding Feature Information 20-1

Prerequisites for Configuring STP 20-1

Restrictions for Configuring STP 20-1

Information About Configuring STP 20-1

STP 20-2

Spanning-Tree Topology and BPDUs 20-2

Bridge ID, Switch Priority, and Extended System ID 20-3

Spanning-Tree Interface States 20-4

Blocking State 20-5

Listening State 20-6

Learning State 20-6

Forwarding State 20-6

Disabled State 20-6

How a Switch or Port Becomes the Root Switch or Root Port 20-7

Spanning Tree and Redundant Connectivity 20-7

Spanning-Tree Address Management 20-8

Accelerated Aging to Retain Connectivity 20-8

Spanning-Tree Modes and Protocols 20-9

Supported Spanning-Tree Instances 20-9

Spanning-Tree Interoperability and Backward Compatibility 20-10

STP and IEEE 802.1Q Trunks 20-10

VLAN-Bridge Spanning Tree 20-10

Default Spanning-Tree Settings 20-11

Disabling Spanning Tree 20-11

Root Switch 20-11

Secondary Root Switch 20-12

Port Priority 20-12

Path Cost 20-13

Spanning-Tree Timers 20-13

Spanning-Tree Configuration Guidelines 20-13

How to Configure STP 20-14

Changing the Spanning-Tree Mode 20-14

Configuring the Root Switch 20-15

Configuring a Secondary Root Switch 20-16

Configuring Port Priority 20-16

Configuring Path Cost 20-16

Configuring Optional STP Parameters 20-17
Monitoring and Maintaining STP 20-17
Additional References 20-18
Related Documents 20-18
Standards 20-18
MIBs 20-18
RFCs 20-18

CHAPTER 21 **Configuring MSTP** 21-1
Finding Feature Information 21-1
Information About Configuring MSTP 21-1
MSTP 21-2
Multiple Spanning-Tree Regions 21-2
IST, CIST, and CST 21-2
Operations Within an MST Region 21-3
Operations Between MST Regions 21-3
IEEE 802.1s Terminology 21-4
Hop Count 21-5
Boundary Ports 21-5
IEEE 802.1s Implementation 21-6
Port Role Naming Change 21-6
Interoperation Between Legacy and Standard Switches 21-6
Detecting Unidirectional Link Failure 21-7
Interoperability with IEEE 802.1D STP 21-8
RSTP 21-8
Port Roles and the Active Topology 21-8
Rapid Convergence 21-9
Synchronization of Port Roles 21-10
Bridge Protocol Data Unit Format and Processing 21-11
Processing Superior BPDU Information 21-12
Processing Inferior BPDU Information 21-12
Topology Changes 21-12
Default MSTP Settings 21-13
MSTP Configuration Guidelines 21-13
Root Switch 21-14
Secondary Root Switch 21-15
Port Priority 21-15
Path Cost 21-15
Link Type to Ensure Rapid Transitions 21-15

Neighbor Type 21-15
Restarting the Protocol Migration Process 21-16
How to Configure MSTP 21-16
Specifying the MST Region Configuration and Enabling MSTP 21-16
Configuring the Root Switch 21-17
Configuring the Optional MSTP Parameters 21-18
Monitoring and Maintaining MSTP 21-20
Configuration Examples for Configuring MSTP 21-20
Configuring the MST Region: Example 21-20
Additional References 21-21
Related Documents 21-21
Standards 21-21
MIBs 21-21
RFCs 21-21

CHAPTER 22 **Configuring Optional Spanning-Tree Features** 22-1
Finding Feature Information 22-1
Prerequisites for the Optional Spanning-Tree Features 22-1
Restrictions for the Optional Spanning-Tree Features 22-1
Information About Configuring the Optional Spanning-Tree Features 22-1
PortFast 22-1
BPDU Guard 22-2
BPDU Filtering 22-3
UplinkFast 22-3
BackboneFast 22-5
EtherChannel Guard 22-7
Root Guard 22-7
Loop Guard 22-8
Default Optional Spanning-Tree Settings 22-9
How to Configure the Optional Spanning-Tree Features 22-9
Enabling Optional SPT Features 22-9
Maintaining and Monitoring Optional Spanning-Tree Features 22-10
Additional References 22-11
Related Documents 22-11
Standards 22-11
MIBs 22-11
RFCs 22-12

CHAPTER 23 **Configuring Resilient Ethernet Protocol** 23-1

Finding Feature Information 23-1

Prerequisites for REP 23-1

Restrictions for REP 23-1

Information About Configuring REP 23-1

REP 23-1

Link Integrity 23-4

Fast Convergence 23-4

VLAN Load Balancing 23-4

Spanning Tree Interaction 23-6

REP Ports 23-6

REP Segments 23-7

Default REP Configuration 23-7

REP Configuration Guidelines 23-7

REP Administrative VLAN 23-8

How to Configure REP 23-9

Configuring the REP Administrative VLAN 23-9

Configuring REP Interfaces 23-9

Setting Manual Preemption for VLAN Load Balancing 23-12

Configuring SNMP Traps for REP 23-12

Monitoring and Maintaining REP 23-12

Configuration Examples for Configuring REP 23-13

Configuring the Administrative VLAN: Example 23-13

Configuring a Primary Edge Port: Examples 23-13

Configuring VLAN Blocking: Example 23-14

Additional References 23-14

Related Documents 23-14

Standards 23-14

MIBs 23-15

RFCs 23-15

CHAPTER 24 **Configuring FlexLinks and the MAC Address-Table Move Update** 24-1

Finding Feature Information 24-1

Restrictions for the FlexLinks and the MAC Address-Table Move Update 24-1

Information About Configuring the FlexLinks and the MAC Address-Table Move Update 24-1

FlexLinks 24-1

VLAN FlexLinks Load Balancing and Support 24-2

FlexLinks Multicast Fast Convergence 24-3

Learning the Other FlexLinks Port as the mrouter Port 24-3
Generating IGMP Reports 24-3
Leaking IGMP Reports 24-4
MAC Address-Table Move Update 24-4
Default Settings for FlexLinks and MAC Address-Table Move Update 24-5
Configuration Guidelines for FlexLinks and MAC Address-Table Move Update 24-6
How to Configure the FlexLinks and MAC Address-Table Move Update 24-6
Configuring FlexLinks 24-6
Configuring a Preemption Scheme for FlexLinks 24-7
Configuring VLAN Load Balancing on FlexLinks 24-7
Configuring the MAC Address-Table Move Update Feature 24-8
Configuring the MAC Address-Table Move Update Messages 24-8
Maintaining and Monitoring the FlexLinks and MAC Address-Table Move Update 24-9
Configuration Examples for the FlexLinks and MAC Address-Table Move Update 24-9
Configuring FlexLinks Port: Examples 24-9
Configuring a Backup Interface: Example 24-11
Configuring a Preemption Scheme: Example 24-11
Configuring VLAN Load Balancing on FlexLinks: Examples 24-12
Configuring MAC Address-Table Move Update: Example 24-13
Additional References 24-13
Related Documents 24-13
Standards 24-13
MIBs 24-14
RFCs 24-14

CHAPTER 25 **Configuring DHCP** 25-1
Finding Feature Information 25-1
Information About Configuring DHCP 25-1
DHCP Snooping 25-1
DHCP Server 25-1
DHCP Relay Agent 25-2
DHCP Snooping 25-2
Option-82 Data Insertion 25-3
Cisco IOS DHCP Server Database 25-6
DHCP Snooping Binding Database 25-6
Default DHCP Snooping Settings 25-7
DHCP Snooping Configuration Guidelines 25-8
DHCP Snooping Binding Database Guidelines 25-9
Packet Forwarding Address 25-9

DHCP Server Port-Based Address Allocation 25-9
How to Configure DHCP 25-10
Configuring the DHCP Relay Agent 25-10
Specifying the Packet Forwarding Address 25-10
Enabling DHCP Snooping and Option 82 25-11
Enabling the DHCP Snooping Binding Database Agent 25-12
Enabling DHCP Server Port-Based Address Allocation 25-13
Preassigning an IP Address 25-13
Monitoring and Maintaining DHCP 25-14
Configuration Examples for Configuring DHCP 25-15
Enabling DHCP Server Port-Based Address Allocation: Examples 25-15
Enabling DHCP Snooping: Example 25-15
Additional References 25-16
Related Documents 25-16
Standards 25-16
MIBs 25-16
RFCs 25-16

CHAPTER 26 **Configuring Dynamic ARP Inspection** 26-1
Finding Feature Information 26-1
Prerequisites for Dynamic ARP Inspection 26-1
Restrictions for Dynamic ARP Inspection 26-1
Information About Dynamic ARP Inspection 26-1
Dynamic ARP Inspection 26-1
Interface Trust States and Network Security 26-3
Rate Limiting of ARP Packets 26-4
Relative Priority of ARP ACLs and DHCP Snooping Entries 26-4
Logging of Dropped Packets 26-4
Default Dynamic ARP Inspection Settings 26-5
Dynamic ARP Inspection Configuration Guidelines 26-5
How to Configure Dynamic ARP Inspection 26-6
Configuring Dynamic ARP Inspection in DHCP Environments 26-6
Configuring ARP ACLs for Non-DHCP Environments 26-7
Limiting the Rate of Incoming ARP Packets 26-9
Performing Validation Checks 26-10
Configuring the Log Buffer 26-11
Monitoring and Maintaining Dynamic ARP Inspection 26-12
Configuration Examples for Dynamic ARP Inspection 26-12

Configuring Dynamic ARP Inspection in DHCP Environments: Example 26-12
Configuring ARP ACLs for Non-DHCP Environments: Example 26-12
Additional References 26-13
Related Documents 26-13
Standards 26-13
MIBs 26-13
RFCs 26-13
Technical Assistance 26-13

CHAPTER 27 **Configuring IP Source Guard** 27-1
Finding Feature Information 27-1
Prerequisites for IP Source Guard 27-1
Restrictions for IP Source Guard 27-1
Information About IP Source Guard 27-1
IP Source Guard 27-1
Source IP Address Filtering 27-2
Source IP and MAC Address Filtering 27-2
IP Source Guard for Static Hosts 27-2
IP Source Guard Configuration Guidelines 27-3
How to Configure IP Source Guard 27-4
Enabling IP Source Guard 27-4
Configuring IP Source Guard for Static Hosts on a Layer 2 Access Port 27-4
Configuring IP Source Guard for Static Hosts on a Private VLAN Host Port 27-5
Monitoring and Maintaining IP Source Guard 27-7
Configuration Examples for IP Source Guard 27-7
Enabling IPSG with Source IP and MAC Filtering: Example 27-7
Disabling IPSG with Static Hosts: Example 27-7
Enabling IPSG for Static Hosts: Examples 27-7
Displaying IP or MAC Binding Entries: Examples 27-8
Enabling IPSG for Static Hosts: Examples 27-9
Additional References 27-10
Related Documents 27-10
Standards 27-11
MIBs 27-11
RFCs 27-11

CHAPTER 28 **Configuring IGMP Snooping and MVR** 28-1
Finding Feature Information 28-1

Restrictions for IGMP Snooping and MVR 28-1
Information About IGMP Snooping and MVR 28-1
IGMP Snooping 28-2
IGMP Versions 28-2
Joining a Multicast Group 28-3
Leaving a Multicast Group 28-5
Immediate Leave 28-5
IGMP Configurable-Leave Timer 28-5
IGMP Report Suppression 28-6
Default IGMP Snooping Configuration 28-6
Snooping Methods 28-6
Multicast Flooding Time After a TCN Event 28-7
Flood Mode for TCN 28-7
Multicast Flooding During a TCN Event 28-7
IGMP Snooping Querier Guidelines 28-7
IGMP Report Suppression 28-8
Multicast VLAN Registration 28-8
MVR in a Multicast Television Application 28-9
Default MVR Settings 28-11
MVR Configuration Guidelines and Limitations 28-11
IGMP Filtering and Throttling 28-12
Default IGMP Filtering and Throttling Configuration 28-12
IGMP Profiles 28-13
IGMP Throttling Action 28-13
How to Configure IGMP Snooping and MVR 28-14
Configuring IGMP Snooping 28-14
Enabling or Disabling IGMP Snooping 28-14
Setting IGMP Snooping Parameters 28-14
Configuring TCN 28-15
Configuring the IGMP Snooping Querier 28-16
Disabling IGMP Report Suppression 28-16
Configuring MVR 28-16
Configuring MVR Global Parameters 28-16
Configuring MVR Interfaces 28-17
Configuring IGMP 28-18
Configuring IGMP Profiles 28-18
Configuring IGMP Interfaces 28-18
Monitoring and Maintaining IGMP Snooping and MVR 28-19
Configuration Examples for IGMP Snooping 28-21

Configuring IGMP Snooping: Example 28-21
Disabling a Multicast Router Port: Example 28-21
Statically Configuring a Host on a Port: Example 28-21
Enabling IGMP Immediate Leave: Example 28-21
Setting the IGMP Snoopng Querier Parameters: Examples 28-21
Enabling MVR: Examples 28-22
Creating an IGMP Profile: Example 28-22
Applying an IGMP Profile: Example 28-23
Limiting IGMP Groups: Example 28-23
Additional References 28-23
Related Documents 28-23
Standards 28-23
MIBs 28-23
RFCs 28-24
Technical Assistance 28-24

CHAPTER 29 **Configuring Port-Based Traffic Control** 29-1
Finding Feature Information 29-1
Restrictions for Port-Based Traffic Control 29-1
Information About Port-Based Traffic Control 29-1
Storm Control 29-1
Default Storm Control Configuration 29-2
Storm Control and Threshold Levels 29-3
Small-Frame Arrival Rate 29-3
Protected Ports 29-3
Protected Port Configuration Guidelines 29-3
Port Blocking 29-4
Port Security 29-4
Secure MAC Addresses 29-4
Security Violations 29-5
Default Port Security Configuration 29-6
Port Security Configuration Guidelines 29-6
Port Security Aging 29-8
Port Security and Private VLANs 29-8
Protocol Storm Protection 29-8
How to Configure Port-Based Traffic Control 29-9
Configuring Storm Control 29-9
Configuring Storm Control and Threshold Levels 29-9
Configuring Small-Frame Arrival Rate 29-10

Configuring Protected Ports 29-10
Configuring Port Blocking 29-11
Blocking Flooded Traffic on an Interface 29-11
Configuring Port Security 29-11
Enabling and Configuring Port Security 29-11
Enabling and Configuring Port Security Aging 29-15
Configuring Protocol Storm Protection 29-15
Enabling Protocol Storm Protection 29-15
Monitoring and Maintaining Port-Based Traffic Control 29-16
Configuration Examples for Port-Based Traffic Control 29-16
Enabling Unicast Storm Control: Example 29-16
Enabling Broadcast Address Storm Control on a Port: Example 29-17
Enabling Small-Frame Arrival Rate: Example 29-17
Configuring a Protected Port: Example 29-17
Blocking Flooding on a Port: Example 29-17
Configuring Port Security: Examples 29-17
Configuring Port Security Aging: Examples 29-18
Configuring Protocol Storm Protection: Example 29-18
Additional References 29-19
Related Documents 29-19
Standards 29-19
MIBs 29-19
RFCs 29-19
Technical Assistance 29-19

CHAPTER 30 **Configuring SPAN and RSPAN** 30-1
Finding Feature Information 30-1
Prerequisites for SPAN and RSPAN 30-1
Restrictions for SPAN and RSPAN 30-1
Information About SPAN and RSPAN 30-1
SPAN and RSPAN 30-1
Local SPAN 30-2
Remote SPAN 30-2
SPAN Sessions 30-3
Monitored Traffic Types for SPAN Sessions 30-4
Source Ports 30-5
Source VLANs 30-6
VLAN Filtering 30-6
Destination Port 30-6

RSPAN VLAN 30-7
SPAN and RSPAN Interaction with Other Features 30-8
Local SPAN Configuration Guidelines 30-9
RSPAN Configuration Guidelines 30-9
Default SPAN and RSPAN Settings 30-10
How to Configure SPAN and RSPAN 30-10
Creating a Local SPAN Session 30-10
Creating a Local SPAN Session and Configuring Incoming Traffic 30-12
Specifying VLANs to Filter 30-13
Configuring a VLAN as an RSPAN VLAN 30-14
Creating an RSPAN Source Session 30-15
Creating an RSPAN Destination Session 30-16
Creating an RSPAN Destination Session and Configuring Incoming Traffic 30-16
Specifying VLANs to Filter 30-17
Monitoring and Maintaining SPAN and RSPAN 30-18
Configuration Examples for SPAN and RSPAN 30-18
Configuring a Local SPAN Session: Example 30-18
Modifying Local SPAN Sessions: Examples 30-18
Configuring an RSPAN: Example 30-19
Configuring a VLAN for a SPAN Session: Example 30-20
Modifying RSPAN Sessions: Examples 30-20
Additional References 30-20
Related Documents 30-20
Standards 30-21
MIBs 30-21
RFCs 30-21

CHAPTER 31 **Configuring LLDP, LLDP-MED, and Wired Location Service** 31-1
Finding Feature Information 31-1
Restrictions for LLDP, LLDP-MED, and Wired Location Service 31-1
Information About LLDP, LLDP-MED, and Wired Location Service 31-1
LLDP-MED 31-2
Wired Location Service 31-3
Default LLDP Configuration 31-4
LLDP, LLDP-MED, and Wired Location Service Configuration Guidelines 31-4
LLDP-MED TLVs 31-5
How to Configure LLDP, LLDP-MED, and Wired Location Service 31-5
Enabling LLDP 31-5
Configuring LLDP Characteristics 31-5

Configuring LLDP-MED TLVs 31-6
Configuring Network-Policy TLV 31-6
Configuring Location TLV and Wired Location Service 31-7
Monitoring and Maintaining LLDP, LLDP-MED, and Wired Location Service 31-8
Configuration Examples for Configuring LLDP, LLDP-MED, and Wired Location Service 31-9
Enabling LLDP: Examples 31-9
Configuring LDP Parameters: Examples 31-9
Configuring TLV: Example 31-9
Configuring Network Policy: Example 31-10
Configuring Voice Application: Example 31-10
Configuring Civic Location Information: Example 31-10
Enabling NMSP: Example 31-10
Additional References 31-11
Related Documents 31-11
Standards 31-11
MIBs 31-11
RFCs 31-11
Technical Assistance 31-11

CHAPTER 32 **Configuring CDP** 32-1
Finding Feature Information 32-1
Information About CDP 32-1
CDP 32-1
Default CDP Configuration 32-2
How to Configure CDP 32-2
Configuring the CDP Parameters 32-2
Disabling CDP 32-3
Monitoring and Maintaining CDP 32-3
Configuration Examples for CDP 32-4
Configuring CDP Parameters: Example 32-4
Enabling CDP: Examples 32-4
Additional References 32-4
Related Documents 32-4
Standards 32-5
MIBs 32-5
RFCs 32-5

CHAPTER 33 **Configuring UDLD** 33-1

Finding Feature Information 33-1
Prerequisites for UDLD 33-1
Restrictions for UDLD 33-1
Information About UDLD 33-1
UDLD 33-1
Modes of Operation 33-2
Methods to Detect Unidirectional Links 33-2
Default UDLD Settings 33-4
How to Configure UDLD 33-4
Enabling UDLD Globally 33-4
Enabling UDLD on an Interface 33-5
Setting and Resetting UDLD Parameters 33-5
Maintaining and Monitoring UDLD 33-6
Additional References 33-6
Related Documents 33-6
Standards 33-6
MIBs 33-6
RFCs 33-6
Technical Assistance 33-7

CHAPTER 34 **Configuring RMON** 34-1

Finding Feature Information 34-1
Prerequisites for RMON 34-1
Restrictions for RMON 34-1
Information About RMON 34-1
RMON 34-1
How to Configure RMON 34-3
Configuring RMON Alarms and Events 34-3
Collecting Group History Statistics on an Interface 34-4
Collecting Group Ethernet Statistics on an Interface 34-4
Monitoring and Maintaining RMON 34-5
Configuration Examples for RMON 34-5
Configuring an RMON Alarm Number: Example 34-5
Creating an RMON Event Number: Example 34-5
Configuring RMON Statistics: Example 34-5
Additional References 34-6
Related Documents 34-6

Standards 34-6
MIBs 34-6
RFCs 34-6
Technical Assistance 34-7

CHAPTER 35 **Configuring System Message Logging** 35-1

Finding Feature Information 35-1
Restrictions for System Message Logging 35-1
Information About System Message Logging 35-1
System Message Logging 35-1
System Log Message Format 35-2
Log Messages 35-2
Message Severity Levels 35-3
Configuring UNIX Syslog Servers 35-3
Logging Messages to a UNIX Syslog Daemon 35-4
Default System Message Logging Configuration 35-5
How to Configure System Message Logging 35-5
Disabling Message Logging 35-5
Setting the Message Display Destination Device 35-6
Synchronizing Log Messages 35-7
Enabling and Disabling Time Stamps on Log Messages 35-8
Enabling and Disabling Sequence Numbers in Log Messages 35-8
Defining the Message Severity Level 35-8
Limiting Syslog Messages Sent to the History Table and to SNMP 35-9
Enabling the Configuration-Change Logger 35-9
Configuring the UNIX System Logging Facility 35-10
Monitoring and Maintaining the System Message Log 35-10
Configuration Examples for the System Message Log 35-10
System Message: Example 35-10
Logging Display: Examples 35-11
Enabling the Logger: Example 35-11
Configuration Log Output: Example 35-11
Additional References 35-12
Related Documents 35-12
Standards 35-12
MIBs 35-12
RFCs 35-12
Technical Assistance 35-13

CHAPTER 36 **Configuring SNMP** 36-1

Finding Feature Information 36-1
Prerequisites for SNMP 36-1
Restrictions for SNMP 36-1
Information About SNMP 36-2
SNMP 36-2
SNMP Versions 36-2
SNMP Manager Functions 36-4
SNMP Agent Functions 36-4
SNMP Community Strings 36-4
Using SNMP to Access MIB Variables 36-5
SNMP Notifications 36-5
SNMP ifIndex MIB Object Values 36-6
Community Strings 36-6
SNMP Notifications 36-6
Default SNMP Settings 36-8
How to Configure SNMP 36-8
Disabling the SNMP Agent 36-8
Configuring Community Strings 36-9
Configuring SNMP Groups and Users 36-10
Configuring SNMP Notifications 36-12
Setting the CPU Threshold Notification Types and Values 36-14
Setting the Agent Contact and Location Information 36-14
Limiting TFTP Servers Used Through SNMP 36-15
Monitoring and Maintaining SNMP 36-15
Configuration Examples for SNMP 36-16
Enabling SNMP Versions: Example 36-16
Permit SNMP Manager Access: Example 36-16
Allow Read-Only Access: Example 36-16
Configure SNMP Traps: Examples 36-16
Associating a User with a Remote Host: Example 36-17
Assigning a String to SNMP: Example 36-17
Additional References 36-17
Related Documents 36-17
Standards 36-17
MIBs 36-18
RFCs 36-18
Technical Assistance 36-18

CHAPTER 37 **Configuring Network Security with ACLs** 37-1

Finding Feature Information 37-1

Restrictions for Network Security with ACLs 37-1

Information About Network Security with ACLs 37-1

ACLs 37-1

Supported ACLs 37-2

Port ACLs 37-2

Handling Fragmented and Unfragmented Traffic 37-3

IPv4 ACLs 37-4

Standard and Extended IPv4 ACLs 37-5

Access List Numbers 37-5

ACL Logging 37-6

Numbered Extended ACL 37-6

Resequencing ACEs in an ACL 37-7

Named Standard and Extended ACLs 37-7

Time Ranges with ACLs 37-8

Comments in ACLs 37-8

IPv4 ACL to a Terminal Line 37-9

IPv4 ACL Application to an Interface Guidelines 37-9

Hardware and Software Handling of IP ACLs 37-10

Troubleshooting ACLs 37-10

Named MAC Extended ACLs 37-11

MAC ACL to a Layer 2 Interface 37-11

How to Configure Network Security with ACLs 37-11

Creating a Numbered Standard ACL 37-11

Creating a Numbered Extended ACL 37-13

Creating Named Standard and Extended ACLs 37-15

Using Time Ranges with ACLs 37-16

Applying an IPv4 ACL to a Terminal Line 37-17

Applying an IPv4 ACL to an Interface 37-17

Creating Named MAC Extended ACLs 37-17

Applying a MAC ACL to a Layer 2 Interface 37-18

Monitoring and Maintaining Network Security with ACLs 37-19

Configuration Examples for Network Security with ACLs 37-19

Creating a Standard ACL: Example 37-19

Creating an Extended ACL: Example 37-19

Configuring Time Ranges: Examples 37-20

Using Named ACLs: Example 37-20

Including Comments in ACLs: Examples 37-21

Applying ACL to a Port: Example 37-21
Applying an ACL to an Interface: Example 37-21
Routed ACLs: Examples 37-22
Configuring Numbered ACLs: Example 37-23
Configuring Extended ACLs: Examples 37-23
Creating Named ACLs: Example 37-24
Applying Time Range to an IP ACL: Example 37-24
Creating Commented IP ACL Entries: Examples 37-25
Configuring ACL Logging: Examples 37-25
Applying a MAC ACL to a Layer 2 Interface: Examples 37-26
Additional References 37-27
Related Documents 37-27
Standards 37-27
MIBs 37-27
RFCs 37-27
Technical Assistance 37-28

CHAPTER 38 **Configuring Standard QoS** 38-1
Finding Feature Information 38-1
Prerequisites for Standard QoS 38-1
Restrictions for Standard QoS 38-1
Information About Standard QoS 38-2
Standard QoS Model 38-4
Standard QoS Configuration Guidelines 38-5
QoS ACL 38-5
QoS on Interfaces 38-5
Policing 38-6
Default Standard QoS Configuration 38-6
Default Ingress Queue Settings 38-7
Default Egress Queue Settings 38-7
Default Mapping Table Settings 38-8
Classification 38-10
Classification Based on QoS ACLs 38-13
Classification Based on Class Maps and Policy Maps 38-13
Policing and Marking 38-14
Policing on Physical Ports 38-15
Policing on SVIs 38-16
Mapping Tables 38-18
Queueing and Scheduling Overview 38-19

Weighted Tail Drop 38-19
SRR Shaping and Sharing 38-20
Queueing and Scheduling on Ingress Queues 38-21
Queueing and Scheduling on Egress Queues 38-22
Packet Modification 38-25
Classification Using Port Trust States 38-26
Trust State on Ports within the QoS Domain 38-26
Configuring a Trusted Boundary to Ensure Port Security 38-26
DSCP Transparency Mode 38-27
DSCP Trust State on a Port Bordering Another QoS Domain 38-27
QoS Policies 38-28
Classifying, Policing, and Marking Traffic on Physical Ports by Using Policy Maps 38-28
Classifying, Policing, and Marking Traffic on SVIs by Using Hierarchical Policy Maps 38-29
DSCP Maps 38-30
DSCP-to-DSCP-Mutation Map 38-30
Ingress Queue Characteristics 38-30
Ingress Priority Queue 38-30
Egress Queue Characteristics 38-31
Egress Queue Configuration Guidelines 38-31
Allocating Buffer Space to and Setting WTD Thresholds for an Egress Queue-Set 38-31
How to Configure Standard QoS 38-32
Enabling QoS Globally 38-32
Enabling VLAN-Based QoS on Physical Ports 38-32
Configuring Classification Using Port Trust States 38-32
Configuring the Trust State on Ports Within the QoS Domain 38-33
Configuring the CoS Value for an Interface 38-33
Configuring a Trusted Boundary to Ensure Port Security 38-34
Enabling DSCP Transparency Mode 38-34
Configuring the DSCP Trust State on a Port Bordering Another QoS Domain 38-35
Configuring a QoS Policy 38-36
Creating IP Standard ACLs 38-36
Creating IP Extended ACLs 38-37
Creating a Layer 2 MAC ACL for Non-IP Traffic 38-37
Creating Class Maps 38-38
Creating Nonhierarchical Policy Maps 38-40
Creating Hierarchical Policy Maps 38-42
Creating Aggregate Policers 38-46
Configuring DSCP Maps 38-47
Configuring the CoS-to-DSCP Map 38-47
Configuring the IP-Precedence-to-DSCP Map 38-48

Configuring the Policed-DSCP Map 38-48
Configuring the DSCP-to-CoS Map 38-48
Configuring the DSCP-to-DSCP-Mutation Map 38-49
Configuring Ingress Queue Characteristics 38-49
Mapping DSCP or CoS Values to an Ingress Queue and Setting WTD Thresholds 38-49
Allocating Buffer Space Between the Ingress Queues 38-50
Allocating Bandwidth Between the Ingress Queues 38-51
Configuring the Ingress Priority Queue 38-51
Configuring Egress Queue Characteristics 38-52
Allocating Buffer Space to and Setting WTD Thresholds for an Egress Queue-Set 38-52
Mapping DSCP or CoS Values to an Egress Queue and to a Threshold ID 38-53
Configuring SRR Shaped Weights on Egress Queues 38-54
Configuring SRR Shared Weights on Egress Queues 38-55
Configuring the Egress Expedite Queue 38-56
Limiting the Bandwidth on an Egress Interface 38-56
Monitoring and Maintaining Standard QoS 38-56
Configuration Examples for Standard QoS 38-57
Configuring the SRR Scheduler: Example 38-57
Configuring DSCP-Trusted State on a Port: Example 38-58
Allowing ACL Permission for IP Traffic: Examples 38-58
Configuring a Class Map: Examples 38-58
Creating a Policy Map: Example 38-59
Creating a Layer 2 MAC ACL: Example 38-59
Creating an Aggregate Policer: Example 38-60
Configuring COS-to-DSCP Map: Example 38-60
Configuring DSCP Maps: Examples 38-61
Configuring an Ingress Queue: Example 38-62
Configuring the Egress Queue: Examples 38-63
Creating a Layer 2 MAC ACL: Example 38-63
Additional References 38-64
Related Documents 38-64
Standards 38-64
MIBs 38-64
RFCs 38-64
Technical Assistance 38-65

CHAPTER 39 **Configuring Auto-QoS** 39-1
Finding Feature Information 39-1
Prerequisites for Auto-QoS 39-1

Restrictions for Auto-QoS 39-1
Information About Auto-QoS 39-2
Auto-QoS 39-2
Generated Auto-QoS Configuration 39-3
Effects of Auto-QoS on the Configuration 39-7
How to Configure Auto-QoS 39-8
Enabling Auto-QoS for VoIP 39-8
Configuring QoS to Prioritize VoIP Traffic 39-9
Monitoring and Maintaining Auto-QoS 39-9
Configuration Examples for Auto-QoS 39-10
Auto-QoS Network: Example 39-10
Enabling Auto-QoS VOIP Trust: Example 39-11
Additional References 39-11
Related Documents 39-11
Standards 39-11
MIBs 39-11
RFCs 39-11
Technical Assistance 39-12
39-12

CHAPTER 40 **Configuring EtherChannels** 40-1
Finding Feature Information 40-1
Restrictions for Configuring EtherChannels 40-1
Information About Configuring EtherChannels 40-1
EtherChannels 40-2
Port-Channel Interfaces 40-3
Port Aggregation Protocol 40-4
PAgP Modes 40-4
PAgP Learn Method and Priority 40-5
PAgP Interaction with Virtual Switches and Dual-Active Detection 40-5
PAgP Interaction with Other Features 40-6
Link Aggregation Control Protocol 40-6
LACP Modes 40-6
LACP Hot-Standby Ports 40-7
LACP Interaction with Other Features 40-7
EtherChannel On Mode 40-8
Load Balancing and Forwarding Methods 40-8
Default EtherChannel Settings 40-10
EtherChannel Configuration Guidelines 40-10

How to Configure EtherChannels 40-11
Configuring Layer 2 EtherChannels 40-11
Configuring EtherChannel Load Balancing 40-14
Configuring the PAgP Learn Method and Priority 40-14
Configuring the LACP Hot-Standby Ports 40-15
Monitoring and Maintaining EtherChannels on the IE 2000 Switch 40-15
Configuration Examples for Configuring EtherChannels 40-16
Configuring EtherChannels: Examples 40-16
Additional References 40-16
Related Documents 40-16
Standards 40-16
MIBs 40-17
RFCs 40-17
Technical Assistance 40-17

CHAPTER 41 **Configuring Static IP Unicast Routing** 41-1
Finding Feature Information 41-1
Restrictions for Static IP Unicast Routing 41-1
Information About Configuring Static IP Unicast Routing 41-1
IP Routing 41-2
Types of Routing 41-2
How to Configure Static IP Unicast Routing 41-3
Steps for Configuring Routing 41-3
Enabling IP Unicast Routing 41-3
Assigning IP Addresses to SVIs 41-3
Configuring Static Unicast Routes 41-4
Monitoring and Maintaining the IP Network 41-4
Additional References for Configuring IP Unicast Routing 41-5
Related Documents 41-5
Standards 41-5
MIBs 41-5
RFCs 41-6
Technical Assistance 41-6

CHAPTER 42 **Configuring IPv6 Host Functions** 42-1
Finding Feature Information 42-1
Prerequisites Configuring IPv6 Host Functions 42-1
Information About Configuring IPv6 Host Functions 42-1

IPv6 42-1
IPv6 Addresses 42-2
Supported IPv6 Host Features 42-2
128-Bit Wide Unicast Addresses 42-3
DNS for IPv6 42-3
ICMPv6 42-3
Neighbor Discovery 42-3
Default Router Preference 42-4
IPv6 Stateless Autoconfiguration and Duplicate Address Detection 42-4
IPv6 Applications 42-4
Dual IPv4 and IPv6 Protocol Stacks 42-4
Static Routes for IPv6 42-5
SNMP and Syslog Over IPv6 42-5
HTTP over IPv6 42-6
Default IPv6 Settings 42-6
How to Configure IPv6 Hosting 42-7
Configuring IPv6 Addressing and Enabling IPv6 Host 42-7
Configuring Default Router Preference 42-8
Configuring IPv6 ICMP Rate Limiting 42-9
Monitoring and Maintaining IPv6 Host Information 42-9
Configuration Examples for IPv6 Host Functions 42-10
Enabling IPv6: Example 42-10
Configuring DRP: Example 42-10
Configuring an IPv6 ICMP Error Message Interval 42-10
Displaying Show Command Output: Examples 42-11
Additional References 42-13
Related Documents 42-13
Standards 42-13
MIBs 42-13
RFCs 42-14
Technical Assistance 42-14

CHAPTER 43 **Configuring Link State Tracking** 43-1
Finding Feature Information 43-1
Restrictions for Configuring Link State Tracking 43-1
Information About Configuring Link State Tracking 43-1
Link State Tracking 43-1
Default Link State Tracking Configuration 43-3
How to Configure Link State Tracking 43-4

Configuring Link State Tracking 43-4
Monitoring and Maintaining Link State Tracking 43-4
Configuration Examples for Configuring Link State Tracking 43-4
Displaying Link State Information: Examples 43-4
Creating a Link State Group: Example 43-5
Additional References 43-5
Related Documents 43-5
Standards 43-5
MIBs 43-6
RFCs 43-6
Technical Assistance 43-6

CHAPTER 44 **Configuring IPv6 MLD Snooping** 44-1
Finding Feature Information 44-1
Prerequisites for Configuring IPv6 MLD Snooping 44-1
Restrictions for Configuring IPv6 MLD Snooping 44-1
Information About Configuring IPv6 MLD Snooping 44-1
IPv6 MLD Snooping 44-1
MLD Messages 44-2
MLD Queries 44-2
Multicast Client Aging Robustness 44-3
Multicast Router Discovery 44-3
MLD Reports 44-3
MLD Done Messages and Immediate-Leave 44-4
Topology Change Notification Processing 44-4
Default MLD Snooping Configuration 44-5
MLD Snooping Configuration Guidelines 44-5
Enabling or Disabling MLD Snooping 44-6
Multicast Router Port 44-6
MLD Immediate Leave 44-6
MLD Snooping Queries 44-6
How to Configure IPv6 MLD Snooping 44-6
Enabling or Disabling MLD Snooping 44-6
Configuring a Static Multicast Group 44-7
Configuring a Multicast Router Port 44-7
Enabling MLD Immediate Leave 44-8
Configuring MLD Snooping Queries 44-8
Disabling MLD Listener Message Suppression 44-9
Monitoring and Maintaining IPv6 MLD Snooping 44-9

Configuration Examples for Configuring IPv6 MLD Snooping 44-10
Statically Configure an IPv6 Multicast Group: Example 44-10
Adding a Multicast Router Port to a VLAN: Example 44-10
Enabling MLD Immediate Leave on a VLAN: Example 44-10
Setting MLD Snooping Global Robustness: Example 44-10
Setting MLD Snooping Last-Listener Query Parameters: Examples 44-10
Additional References 44-12
Related Documents 44-12
Standards 44-12
MIBs 44-12
RFCs 44-12
Technical Assistance 44-12

CHAPTER 45 **Configuring Cisco IOS IP SLAs Operations** 45-1
Finding Feature Information 45-1
Prerequisites for Configuring Cisco IOS IP SLAs Operations 45-1
Restrictions for Configuring Cisco IOS IP SLAs Operations 45-1
Information About Configuring Cisco IOS IP SLAs Operations 45-1
Cisco IOS IP SLAs 45-2
Cisco IOS IP SLAs to Measure Network Performance 45-3
IP SLAs Responder and IP SLAs Control Protocol 45-3
Response Time Computation for IP SLAs 45-4
IP SLAs Operation Scheduling 45-4
IP SLAs Operation Threshold Monitoring 45-5
IP Service Levels by Using the UDP Jitter Operation 45-5
IP Service Levels by Using the ICMP Echo Operation 45-6
How to Configure Cisco IOS IP SLAs Operations 45-6
Configuring the IP SLAs Responder 45-7
Configuring UDP Jitter Operation 45-7
Analyzing IP Service Levels by Using the ICMP Echo Operation 45-9
Monitoring and Maintaining Cisco IP SLAs Operations 45-10
Configuration Examples for Configuring Cisco IP SLAs Operations 45-11
Configuring an ICMP Echo IP SLAs Operation: Example 45-11
Sample Output for Show IP SLA Command: Example 45-12
Configuring a Responder UDP Jitter IP SLAs Operation: Example 45-12
Configuring a UDP Jitter IP SLAs Operation: Example 45-12
Additional References 45-13
Related Documents 45-13
Standards 45-13

MIBs 45-14
RFCs 45-14
Technical Assistance 45-14

CHAPTER 46 **Troubleshooting** 46-1

Finding Feature Information 46-1
Information for Troubleshooting 46-1
Autonegotiation Mismatches Prevention 46-1
SFP Module Security and Identification 46-2
Ping 46-2
Layer 2 Traceroute 46-3
Layer 2 Traceroute Usage Guidelines 46-3
IP Traceroute 46-4
TDR 46-4
Crashinfo Files 46-5
Basic crashinfo Files 46-5
Extended crashinfo Files 46-5
CPU Utilization 46-6
Problem and Cause for High CPU Utilization 46-6
How to Troubleshoot 46-7
Recovering from Software Failures 46-7
Recovering from a Lost or Forgotten Password 46-8
Recovering from Lost Cluster Member Connectivity 46-9
Executing Ping 46-9
Executing IP Traceroute 46-10
Running TDR and Displaying the Results 46-11
Enabling Debugging on a Specific Feature 46-12
Enabling All-System Diagnostics 46-12
Redirecting Debug and Error Message Output 46-13
Monitoring Information 46-13
Physical Path 46-13
SFP Module Status 46-13
Troubleshooting Examples 46-14
show platform forward Command 46-14
Additional References 46-16
Related Documents 46-16
Standards 46-16
MIBs 46-16
RFCs 46-17

Technical Assistance 46-17

APPENDIX A **Working with the Cisco IOS File System, Configuration Files, and Software Images A-1**

Working with the Flash File System A-1
Displaying Available File Systems A-1
Detecting an Unsupported SD Flash Memory Card A-2
SD Flash Memory Card LED A-3
Setting the Default File System A-3
Displaying Information About Files on a File System A-4
Changing Directories and Displaying the Working Directory A-5
Creating and Removing Directories A-5
Copying Files A-6
Deleting Files A-6
Creating, Displaying, and Extracting tar Files A-7
Creating a tar File A-7
Displaying the Contents of a tar File A-7
Extracting a tar File A-8
Displaying the Contents of a File A-9
Working with Configuration Files A-9
Guidelines for Creating and Using Configuration Files A-9
Configuration File Types and Location A-10
Creating a Configuration File By Using a Text Editor A-10
Copying Configuration Files By Using TFTP A-11
Preparing to Download or Upload a Configuration File By Using TFTP A-11
Downloading the Configuration File By Using TFTP A-11
Uploading the Configuration File By Using TFTP A-12
Copying Configuration Files By Using FTP A-13
Preparing to Download or Upload a Configuration File By Using FTP A-13
Downloading a Configuration File By Using FTP A-14
Uploading a Configuration File By Using FTP A-15
Copying Configuration Files By Using RCP A-16
Preparing to Download or Upload a Configuration File By Using RCP A-16
Downloading a Configuration File By Using RCP A-17
Uploading a Configuration File By Using RCP A-18
Clearing Configuration Information A-19
Clearing the Startup Configuration File A-19
Deleting a Stored Configuration File A-19
Replacing and Rolling Back Configurations A-19
Understanding Configuration Replacement and Rollback A-19

Configuration Guidelines A-20
Configuring the Configuration Archive A-21
Performing a Configuration Replacement or Rollback Operation A-21
Working with Software Images A-22
Image Location on the Switch A-23
tar File Format of Images on a Server or Cisco.com A-23
Copying Image Files By Using TFTP A-24
Preparing to Download or Upload an Image File By Using TFTP A-25
Downloading an Image File By Using TFTP A-25
Uploading an Image File By Using TFTP A-27
Copying Image Files By Using FTP A-27
Preparing to Download or Upload an Image File By Using FTP A-28
Downloading an Image File By Using FTP A-29
Uploading an Image File By Using FTP A-30
Copying Image Files By Using RCP A-31
Preparing to Download or Upload an Image File By Using RCP A-32
Downloading an Image File By Using RCP A-33
Uploading an Image File By Using RCP A-34

INDEX

Preface

Audience

This guide is for the networking professional managing your switch. Before using this guide, you should have experience working with the Cisco IOS software and be familiar with the concepts and terminology of Ethernet and local area networking.

Purpose

This guide provides the information that you need to configure Cisco IOS software features on your switch.

This guide provides procedures for using the commands that have been created or changed for use with the switch. It does not provide detailed information about these commands. For detailed information about these commands, see the *Cisco IE 2000 Switch Command Reference* for this release.

For information about the standard Cisco IOS commands, see the Cisco IOS 15.0 documentation set available from the Cisco.com home page.

This guide does not provide detailed information on the graphical user interfaces (GUIs) for the embedded Device Manager. However, the concepts in this guide are applicable to the GUI user. For information about Device Manager, see the switch online help.

For documentation updates, see the release notes for this release.

Conventions

This publication uses these conventions to convey instructions and information:

Command descriptions use these conventions:

- Commands and keywords are in **boldface** text.
- Arguments for which you supply values are in *italic*.
- Square brackets ([ ]) mean optional elements.
- Braces ({ }) group required choices, and vertical bars (|) separate the alternative elements.
- Braces and vertical bars within square brackets ([{ | }]) mean a required choice within an optional element.

Interactive examples use these conventions:

- Terminal sessions and system displays are in `screen` font.
- Information you enter is in **`boldface screen`** font.
- Nonprinting characters, such as passwords or tabs, are in angle brackets (< >).

Notes, cautions, and timesavers use these conventions and symbols:

Note Means *reader take note*. Notes contain helpful suggestions or references to materials not contained in this manual.

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Related Publications

These documents provide complete information about the switch and are available from this Cisco.com site:

http://www.cisco.com/go/ie2000_docs

Note Before installing, configuring, or upgrading the switch, see these documents:

- For initial configuration information, see the "Using Express Setup" section in the getting started guide or the "Configuring the Switch with the CLI-Based Setup Program" appendix in the hardware installation guide.
- For Device Manager requirements, see the "System Requirements" section in the release notes (not orderable but available on Cisco.com).
- For upgrading information, see the "Downloading Software" section in the release notes.

See these documents for other information about the switch:

- *Release Notes for the Cisco IE 2000 Switch*
- *Cisco IE 2000 Switch Software Configuration Guide*
- *Cisco IE 2000 Switch Command Reference*
- *Cisco IE 2000 Switch System Message Guide*
- *Cisco IE 2000 Switch Hardware Installation Guide*
- *Cisco IE 2000 Switch Getting Started Guide*
- *Regulatory Compliance and Safety Information for the Cisco IE 2000 Switch*
- *Cisco Small Form-Factor Pluggable Modules Installation Notes*
- Device Manager online help (available on the switch)
- For more information about the Network Admission Control (NAC) features, see the *Network Admission Control Software Configuration Guide.*
- Compatibility matrix documents are available from this Cisco.com site:

http://www.cisco.com/en/US/products/hw/modules/ps5455/products_device_support_tables_list.html

- Cisco Gigabit Ethernet Transceiver Modules Compatibility Matrix

Obtaining Documentation, Obtaining Support, and Security Guidelines

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.

CHAPTER 1

Configuration Overview

Features

Your switch uses the Cisco IOS software licensing (CISL) architecture to support a single universal cryptographic image (supports encryption). This image implements the LAN Base or LAN Lite features depending on your switch model:r

- The LAN Base image provides quality of service (QoS), port security, 1588v2 PTP, and static routing features.
- The LAN Lite image provides reduced Layer 2 functionality without the loss of critical security features such as SSH and SNMPv3.

Feature Software Licensing

A feature license is supported on a single universal image that implements the LAN Base or LAN Lite features depending on your software license:

- The LAN Base features include quality of service (QoS), port security, PTP, and static routing.
- The LAN Lite features provide Layer 2 functionality without losing critical security features such as SSH and SNMPv3.

Cryptographic functionality is included on the universal image.

These guidelines can help you determine what image is running on your switch:

- Enter the **show version** privileged EXEC command. For example, IE-2000-8TC-G-E runs the LAN Base image by default and the IE-2000-4T-G-L runs the LAN Lite image by default.
- Enter the **show license** privileged EXEC command, to see which is the active image:

```
Switch# show license
Index 1 Feature: lanbase
        Period left: Life time
        License Type: Permanent
        License State: Active, In Use
        License Priority: Medium
        License Count: Non-Counted

Index 2 Feature: lanlite
        Period left: 0  minute  0  second
```

Ease-of-Deployment and Ease-of-Use Features

- Express Setup for quickly configuring a switch for the first time with basic IP information, contact information, switch and Telnet passwords, and Simple Network Management Protocol (SNMP) information through a browser-based program. For more information about Express Setup, see the getting started guide.
- User-defined and Cisco-default Smartports macros for creating custom switch configurations for simplified deployment across the network.
- A removable SD flash card that stores the Cisco IOS software image and configuration files for the switch. You can replace and upgrade the switch without reconfiguring the software features.
- An embedded Device Manager GUI for configuring and monitoring a single switch through a web browser. For information about launching Device Manager, see the getting started guide. For more information about Device Manager, see the switch online help.

Performance Features

- Autosensing of port speed and autonegotiation of duplex mode on all switch ports for optimizing bandwidth
- Automatic medium-dependent interface crossover (auto-MDIX) capability on 10/100 and 10/100/1000 Mb/s interfaces and on 10/100/1000 BASE-TX SFP module interfaces that enables the interface to automatically detect the required cable connection type (straight-through or crossover) and to configure the connection appropriately
- Support for up to 1546 bytes routed frames, up to 9000 bytes for frames that are bridged in hardware, and up to 2000 bytes for frames that are bridged by software
- IEEE 802.3x flow control on all ports (the switch does not send pause frames)
- Support for up to 6 EtherChannel groups
- Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol (LACP) for automatic creation of EtherChannel links
- Per-port storm control for preventing broadcast, multicast, and unicast storms
- Port blocking on forwarding unknown Layer 2 unknown unicast, multicast, and bridged broadcast traffic
- Cisco Group Management Protocol (CGMP) server support and Internet Group Management Protocol (IGMP) snooping for IGMP Versions 1, 2, and 3:
  - (For CGMP devices) CGMP for limiting multicast traffic to specified end stations and reducing overall network traffic
  - (For IGMP devices) IGMP snooping for forwarding multimedia and multicast traffic
- IGMP report suppression for sending only one IGMP report per multicast router query to the multicast devices (supported only for IGMPv1 or IGMPv2 queries)
- IGMP snooping querier support to configure switch to generate periodic IGMP general query messages
- IGMP helper to allow the switch to forward a host request to join a multicast stream to a specific IP destination address
- IGMP filtering for controlling the set of multicast groups to which hosts on a switch port can belong

- IGMP throttling for configuring the action when the maximum number of entries is in the IGMP forwarding table
- IGMP leave timer for configuring the leave latency for the network
- Switch Database Management (SDM) templates for allocating system resources to maximize support for user-selected features
- Cisco IOS IP Service Level Agreements (SLAs), a part of Cisco IOS software that uses active traffic monitoring for measuring network performance
- Configurable small-frame arrival threshold to prevent storm control when small frames (64 bytes or less) arrive on an interface at a specified rate (the threshold)
- FlexLink Multicast Fast Convergence to reduce the multicast traffic convergence time after a FlexLink failure
- RADIUS server load balancing to allow access and authentication requests to be distributed evenly across a server group
- Support for QoS marking of CPU-generated traffic and queue CPU-generated traffic on the egress network ports

Management Options

- An embedded Device Manager—Device Manager is a GUI application that is integrated in the software image. You use it to configure and to monitor a single switch. For information about launching Device Manager, see the getting started guide. For more information about Device Manager, see the switch online help.
- Network Assistant—Network Assistant is a network management application that can be downloaded from Cisco.com. You use it to manage a single switch, a cluster of switches, or a community of devices. For more information about Network Assistant, see *Getting Started with Cisco Network Assistant*, available on Cisco.com.
- CLI—The Cisco IOS software supports desktop- and multilayer-switching features. You can access the CLI either by connecting your management station directly to the switch console port or by using Telnet from a remote management station. For more information about the CLI, see Chapter 2, “Using the Command-Line Interface.”
- SNMP—SNMP management applications such as CiscoWorks2000 LAN Management Suite (LMS) and HP OpenView. You can manage from an SNMP-compatible management station that is running platforms such as HP OpenView or SunNet Manager. The switch supports a comprehensive set of MIB extensions and four remote monitoring (RMON) groups. For more information about using SNMP, see Chapter 36, “Configuring SNMP.”
- Cisco IOS Configuration Engine (previously known as the Cisco IOS CNS agent)—Configuration service automates the deployment and management of network devices and services. You can automate initial configurations and configuration updates by generating switch-specific configuration changes, sending them to the switch, executing the configuration change, and logging the results.

  For more information about CNS, see Chapter 5, “Configuring Cisco IOS Configuration Engine.”

Industrial Application

- CIP—Common Industrial Protocol (CIP) is a peer-to-peer application protocol that provides application level connections between the switch and industrial devices such as I/O controllers, sensors, relays, and so forth.You can manage the switch using CIP-based management tools, such as RSLogix. For more information about the CIP commands that the switch supports, see the command reference.
- Profinet Version 2—Support for PROFINET IO, a modular communication framework for distributed automation applications. The switch provides a PROFINET management connection to the I/O controllers.

Manageability Features

- CNS embedded agents for automating switch management, configuration storage, and delivery.
- DHCP for automating configuration of switch information (such as IP address, default gateway, hostname, and Domain Name System [DNS] and TFTP server names).
- DHCP relay for forwarding User Datagram Protocol (UDP) broadcasts, including IP address requests, from DHCP clients.
- DHCP server for automatic assignment of IP addresses and other DHCP options to IP hosts.
- DHCP-based autoconfiguration and image update to download a specified configuration of a new image to a large number of switches.
- DHCPv6 bulk-lease query to support new bulk lease query type (as defined in RFC5460).
- DHCPv6 Relay Source Configuration feature to configure a source address for DHCPv6 relay agent.
- DHCP server port-based address allocation for the preassignment of an IP address to a switch port.
- Directed unicast requests to a DNS server for identifying a switch through its IP address and its corresponding hostname and to a TFTP server for administering software upgrades from a TFTP server.
- Address Resolution Protocol (ARP) for identifying a switch through its IP address and its corresponding MAC address.
- Unicast MAC address filtering to drop packets with specific source or destination MAC addresses.
- Configurable MAC address scaling that allows disabling MAC address learning on a VLAN to limit the size of the MAC address table.
- Cisco Discovery Protocol (CDP) Versions 1 and 2 for network topology discovery and mapping between the switch and other Cisco devices on the network.
- Link Layer Discovery Protocol (LLDP) and LLDP Media Endpoint Discovery (LLDP-MED) for interoperability with third-party IP phones.
- LLDP media extensions (LLDP-MED) location TLV that provides location information from the switch to the endpoint device.
- Network Time Protocol (NTP) for providing a consistent time stamp to all switches from an external source.
- Network Time Protocol version 4 (NTPv4) to support both IPv4 and IPv6 and compatibility with NTPv3.
- Precision Time Protocol (PTP) as defined in the IEEE 1588 standard to synchronize with nanosecond accuracy the real-time clocks of the devices in a network.

  - PTP enhancement to support PTP messages on the expansion module ports.
- Cisco IOS File System (IFS) for providing a single interface to all file systems that the switch uses.
- Support for the SSM PIM protocol to optimize multicast applications, such as video.
- Configuration logging to log and to view changes to the switch configuration.
- Unique device identifier to provide product identification information through a **show inventory** user EXEC command display.
- In-band management access through Device Manager over a Netscape Navigator or Microsoft Internet Explorer browser session.
- In-band management access for up to 16 simultaneous Telnet connections for multiple CLI-based sessions over the network.
- In-band management access for up to five simultaneous, encrypted Secure Shell (SSH) connections for multiple CLI-based sessions over the network.
- In-band management access through SNMP Versions 1, 2c, and 3 get and set requests.
- Out-of-band management access through the switch console port to a directly attached terminal or to a remote terminal through a serial connection or a modem.
- Secure Copy Protocol (SCP) feature to provide a secure and authenticated method for copying switch configuration or switch image files (requires the cryptographic version of the software).
- Configuration replacement and rollback to replace the running configuration on a switch with any saved Cisco IOS configuration file.
- The HTTP client in Cisco IOS can send requests to both IPv4 and IPv6 HTTP server, and the HTTP server in Cisco IOS can service HTTP requests from both IPv4 and IPv6 HTTP clients.
- Simple Network and Management Protocol (SNMP) can be configured over IPv6 transport so that an IPv6 host can send SNMP queries and receive SNMP notifications from a device running IPv6.
- IPv6 stateless autoconfiguration to manage link, subnet, and site addressing changes, such as management of host and mobile IP addresses.
- Disabling MAC address learning on a VLAN.
- DHCP server port-based address allocation for the preassignment of an IP address to a switch port.
- CPU utilization threshold trap monitors CPU utilization.
- LLDP-MED network-policy profile time, length, value (TLV) for creating a profile for voice and voice-signaling by specifying the values for VLAN, class of service (CoS), differentiated services code point (DSCP), and tagging mode.
- Support for including a hostname in the option 12 field of DHCPDISCOVER packets. This provides identical configuration files to be sent by using the DHCP protocol.
- DHCP Snooping enhancement to support the selection of a fixed string-based format for the circuit-id sub-option of the Option 82 DHCP field.
- Support for PROFINET IO, a modular communication framework for distributed automation applications. The switch provides a PROFINET management connection to the I/O controllers.

Availability and Redundancy Features

- UniDirectional Link Detection (UDLD) and aggressive UDLD for detecting and disabling unidirectional links on fiber-optic interfaces caused by incorrect fiber-optic wiring or port faults

- IEEE 802.1D Spanning Tree Protocol (STP) for redundant backbone connections and loop-free networks. STP has these features:
  - Up to 128 spanning-tree instances supported
  - Per-VLAN spanning-tree plus (PVST+) for load balancing across VLANs
  - Rapid PVST+ for load balancing across VLANs and providing rapid convergence of spanning-tree instances
- IEEE 802.1s Multiple Spanning Tree Protocol (MSTP) for grouping VLANs into a spanning-tree instance and for providing multiple forwarding paths for data traffic and load balancing and rapid per-VLAN Spanning-Tree plus (rapid-PVST+) based on the IEEE 802.1w Rapid Spanning Tree Protocol (RSTP) for rapid convergence of the spanning tree by immediately changing root and designated ports to the forwarding state
- Optional spanning-tree features available in PVST+, rapid-PVST+, and MSTP mode:
  - Port Fast for eliminating the forwarding delay by enabling a port to immediately change from the blocking state to the forwarding state
  - BPDU guard for shutting down Port Fast-enabled ports that receive bridge protocol data units (BPDUs)
  - BPDU filtering for preventing a Port Fast-enabled port from sending or receiving BPDUs
  - Root guard for preventing switches outside the network core from becoming the spanning-tree root
  - Loop guard for preventing alternate or root ports from becoming designated ports because of a failure that leads to a unidirectional link
- FlexLink Layer 2 interfaces to back up one another as an alternative to STP for basic link redundancy (requires the LAN Base image)
- Link-state tracking to mirror the state of the ports that carry upstream traffic from connected hosts and servers, and to allow the failover of the server traffic to an operational link on another Cisco Ethernet switch.

VLAN Features

- Support for up to 255 VLANs for assigning users to VLANs associated with appropriate network resources, traffic patterns, and bandwidth.
- Support for VLAN IDs in the 1 to 4096 range as allowed by the IEEE 802.1Q standard.
- VLAN Query Protocol (VQP) for dynamic VLAN membership.
- IEEE 802.1Q trunking encapsulation on all ports for network moves, adds, and changes; management and control of broadcast and multicast traffic; and network security by establishing VLAN groups for high-security users and network resources.
- Dynamic Trunking Protocol (DTP) for negotiating trunking on a link between two devices and for negotiating the type of trunking encapsulation (IEEE 802.1Q) to be used.
- VLAN Trunking Protocol (VTP) and VTP pruning for reducing network traffic by restricting flooded traffic to links destined for stations receiving the traffic.
- Voice VLAN for creating subnets for voice traffic from Cisco IP phones.

- VLAN 1 minimization for reducing the risk of spanning-tree loops or storms by allowing VLAN 1 to be disabled on any individual VLAN trunk link. With this feature enabled, no user traffic is sent or received on the trunk. The switch CPU continues to send and receive control protocol frames.
- VLAN FlexLink load balancing to provide Layer 2 redundancy without requiring Spanning Tree Protocol (STP). A pair of interfaces configured as primary and backup links can load balance traffic based on VLAN.
- Support for 802.1x authentication with restricted VLANs (also known as *authentication failed VLANs*).
- Support for VTP version 3 that includes support for configuring extended range VLANs (VLANs 1006 to 4096) in any VTP mode, enhanced authentication (hidden or secret passwords), propagation of other databases in addition to VTP, VTP primary and secondary servers, and the option to turn VTP on or off by port.

Security Features

- IP Service Level Agreements (IP SLAs) support to measure network performance by using active traffic monitoring
- IP SLAs EOT to use the output from IP SLAs tracking operations triggered by an action such as latency, jitter, or packet loss for a standby router failover takeover (requires the LAN Base image)
- Web authentication to allow a supplicant (client) that does not support IEEE 802.1x functionality to be authenticated using a web browser
- Local web authentication banner so that a custom banner or an image file can be displayed at a web authentication login screen
- MAC authentication bypass (MAB) aging timer to detect inactive hosts that have authenticated after they have authenticated by using MAB
- Password-protected access (read-only and read-write access) to management interfaces (Device Manager, Network Assistant, and the CLI) for protection against unauthorized configuration changes
- Multilevel security for a choice of security level, notification, and resulting actions
- Static MAC addressing for ensuring security
- Protected port option for restricting the forwarding of traffic to designated ports on the same switch
- Port security option for limiting and identifying MAC addresses of the stations allowed to access the port
- VLAN-aware port security option to shut down the VLAN on the port when a violation occurs, instead of shutting down the entire port
- Port security aging to set the aging time for secure addresses on a port
- Protocol storm protection to control the rate of incoming protocol traffic to a switch by dropping packets that exceed a specified ingress rate
- BPDU guard for shutting down a Port Fast-configured port when an invalid configuration occurs
- Standard and extended IP access control lists (ACLs) for defining security policies in both directions on routed interfaces (router ACLs) and VLANs and inbound on Layer 2 interfaces (port ACLs)
- Extended MAC access control lists for defining security policies in the inbound direction on Layer 2 interfaces
- Source and destination MAC-based ACLs for filtering non-IP traffic

- DHCP snooping to filter untrusted DHCP messages between untrusted hosts and DHCP servers
- IP source guard to restrict traffic on nonrouted interfaces by filtering traffic based on the DHCP snooping database and IP source bindings
- Dynamic ARP inspection to prevent malicious attacks on the switch by not relaying invalid ARP requests and responses to other ports in the same VLAN
- Layer 2 protocol tunneling bypass feature to provide interoperability with third-party vendors
- IEEE 802.1x port-based authentication to prevent unauthorized devices (clients) from gaining access to the network. These features are supported:
  - Multidomain authentication (MDA) to allow both a data device and a voice device, such as an IP phone (Cisco or non-Cisco), to independently authenticate on the same IEEE 802.1x-enabled switch port
  - Dynamic voice virtual LAN (VLAN) for MDA to allow a dynamic voice VLAN on an MDA-enabled port
  - VLAN assignment for restricting 802.1x-authenticated users to a specified VLAN
  - Port security for controlling access to 802.1x ports
  - Voice VLAN to permit a Cisco IP Phone to access the voice VLAN regardless of the authorized or unauthorized state of the port
  - IP phone detection enhancement to detect and recognize a Cisco IP phone
  - Guest VLAN to provide limited services to non-802.1x-compliant users
  - Restricted VLAN to provide limited services to users who are 802.1x compliant, but do not have the credentials to authenticate via the standard 802.1x processes
  - 802.1x accounting to track network usage
  - 802.1x with wake-on-LAN to allow dormant PCs to be powered on based on the receipt of a specific Ethernet frame
  - 802.1x readiness check to determine the readiness of connected end hosts before configuring IEEE 802.1x on the switch
  - Voice-aware 802.1x security to apply traffic violation actions only on the VLAN on which a security violation occurs
  - MAC authentication bypass to authorize clients based on the client MAC address
  - Network Edge Access Topology (NEAT) with 802.1X switch supplicant, host authorization with CISP, and auto enablement to authenticate a switch outside a wiring closet as a supplicant to another switch
  - IEEE 802.1x with open access to allow a host to access the network before being authenticated
  - IEEE 802.1x authentication with downloadable ACLs and redirect URLs to allow per-user ACL downloads from a Cisco Secure ACS server to an authenticated switch
  - Flexible-authentication sequencing to configure the order of the authentication methods that a port tries when authenticating a new host
  - Multiple-user authentication to allow more than one host to authenticate on an 802.1x-enabled port
- Network Admission Control (NAC) features:
  - NAC Layer 2 802.1x validation of the antivirus condition or *posture* of endpoint systems or clients before granting the devices network access

For information about configuring NAC Layer 2 802.1x validation, see the "Configuring NAC Layer 2 802.1x Validation" section on page 13-46

  - NAC Layer 2 IP validation of the posture of endpoint systems or clients before granting the devices network access

    For information about configuring NAC Layer 2 IP validation, see the *Network Admission Control Software Configuration Guide*

  - IEEE 802.1x inaccessible authentication bypass

    For information about configuring this feature, see the "Configuring Inaccessible Authentication Bypass" section on page 13-44

  - Authentication, authorization, and accounting (AAA) down policy for a NAC Layer 2 IP validation of a host if the AAA server is not available when the posture validation occurs

    For information about this feature, see the *Network Admission Control Software Configuration Guide.*

- TACACS+, a proprietary feature for managing network security through a TACACS server
- RADIUS for verifying the identity of, granting access to, and tracking the actions of remote users through AAA services
- Enhancements to RADIUS, TACACS+, and SSH to function over IPv6
- Kerberos security system to authenticate requests for network resources by using a trusted third party (requires the cryptographic versions of the software)
- Secure Socket Layer (SSL) Version 3.0 support for the HTTP 1.1 server authentication, encryption, and message integrity and HTTP client authentication to allow secure HTTP communications (requires the cryptographic version of the software)
- Voice-aware IEEE 802.1x and MAC authentication bypass (MAB) security violation to shut down only the data VLAN on a port when a security violation occurs
- Support for IP source guard on static hosts
- RADIUS change of authorization (CoA) to change the attributes of a certain session after it is authenticated. When there is a change in policy for a user or user group in AAA, administrators can send the RADIUS CoA packets from the AAA server, such as Cisco Secure ACS to reinitialize authentication, and apply to the new policies.
- IEEE 802.1x User Distribution to allow deployments with multiple VLANs (for a group of users) to improve scalability of the network by load balancing users across different VLANs. Authorized users are assigned to the least populated VLAN in the group, assigned by RADIUS server.
- Support for critical VLAN with multiple-host authentication so that when a port is configured for multi-authentication, and an AAA server becomes unreachable, the port is placed in a critical VLAN in order to still permit access to critical resources
- Customizable web authentication enhancement to allow the creation of user-defined *login*, *success*, *failure* and *expire* web pages for local web authentication
- Support for Network Edge Access Topology (NEAT) to change the port host mode and to apply a standard port configuration on the authenticator switch port
- VLAN-ID based MAC authentication to use the combined VLAN and MAC address information for user authentication to prevent network access from unauthorized VLANs
- MAC move to allow hosts (including the hosts connected behind an IP phone) to move across ports within the same switch without any restrictions to enable mobility. With MAC move, the switch treats the reappearance of the same MAC address on another port in the same way as a completely new MAC address.

- Support for 3DES and AES with version 3 of the Simple Network Management Protocol (SNMPv3). This release adds support for the 168-bit Triple Data Encryption Standard (3DES) and the 128-bit, 192-bit, and 256-bit Advanced Encryption Standard (AES) encryption algorithms to SNMPv3.

QoS and CoS Features

Note These features require the LAN Base image.

- Automatic QoS (auto-QoS) to simplify the deployment of existing QoS features by classifying traffic and configuring egress queues
- Automatic quality of service (QoS) Voice over IP (VoIP) enhancement for port-based trust of DSCP and priority queuing for egress traffic
- Classification
  - IP type-of-service/Differentiated Services Code Point (IP ToS/DSCP) and IEEE 802.1p CoS marking priorities on a per-port basis for protecting the performance of mission-critical applications
  - IP ToS/DSCP and IEEE 802.1p CoS marking based on flow-based packet classification (classification based on information in the MAC, IP, and TCP/UDP headers) for high-performance quality of service at the network edge, allowing for differentiated service levels for different types of network traffic and for prioritizing mission-critical traffic in the network
  - Trusted port states (CoS, DSCP, and IP precedence) within a QoS domain and with a port bordering another QoS domain
  - Trusted boundary for detecting the presence of a Cisco IP Phone, trusting the CoS value received, and ensuring port security
- Policing
  - Traffic-policing policies on the switch port for managing how much of the port bandwidth should be allocated to a specific traffic flow.
  - If you configure multiple class maps for a hierarchical policy map, each class map can be associated with its own port-level (second-level) policy map. Each second-level policy map can have a different policer.
  - Aggregate policing for policing traffic flows in aggregate to restrict specific applications or traffic flows to metered, predefined rates.
- Out-of-profile
  - Out-of-profile markdown for packets that exceed bandwidth utilization limits
- Ingress queueing and scheduling
  - Two configurable ingress queues for user traffic (one queue can be the priority queue)
  - Weighted tail drop (WTD) as the congestion-avoidance mechanism for managing the queue lengths and providing drop precedences for different traffic classifications
  - Shaped round robin (SRR) as the scheduling service for specifying the rate at which packets are sent to the ring (sharing is the only supported mode on ingress queues)
- Egress queues and scheduling
  - Four egress queues per port.

  - WTD as the congestion-avoidance mechanism for managing the queue lengths and providing drop precedences for different traffic classifications.
  - SRR as the scheduling service for specifying the rate at which packets are dequeued to the egress interface (shaping or sharing is supported on egress queues). Shaped egress queues are guaranteed but limited to using a share of port bandwidth. Shared egress queues are also guaranteed a configured share of bandwidth, but can use more than the guarantee if other queues become empty and do not use their share of the bandwidth.

Monitoring Features

- EOT and IP SLAs EOT static route support identify when a preconfigured static route or a DHCP route goes down
- MAC address notification traps and RADIUS accounting for tracking users on a network by storing the MAC addresses that the switch has learned or removed
- Switched Port Analyzer (SPAN) and Remote SPAN (RSPAN) for traffic monitoring on any port or VLAN (RSPAN requires LAN Base image)
- SPAN and RSPAN support of Intrusion Detection Systems (IDS) to monitor, repel, and report network security violations (RSPAN requires LAN Base image)
- Four groups (history, statistics, alarms, and events) of embedded RMON agents for network monitoring and traffic analysis
- Syslog facility for logging system messages about authentication or authorization errors, resource issues, and time-out events
- Layer 2 traceroute to identify the physical path that a packet takes from a source device to a destination device
- Time Domain Reflector (TDR) to diagnose and resolve cabling problems on 10/100 and 10/100/1000 copper Ethernet ports
- SFP module diagnostic management interface to monitor physical or operational status of an SFP module
- Facilities for processing alarms related to temperature, power-supply conditions, and the status of the Ethernet ports
- Alarm relay contacts that can be used for an external relay system
- Digital optical monitoring (DOM) to check status of X2 small form-factor pluggable (SFP) modules

Default Settings After Initial Switch Configuration

The switch is designed for plug-and-play operation, requiring only that you assign basic IP information to the switch and connect it to the other devices in your network. If you have specific network needs, you can change the interface-specific and system-wide settings.

Note For information about assigning an IP address by using the browser-based Express Setup program, see the getting started guide. For information about assigning an IP address by using the CLI-based setup program, see the hardware installation guide.

If you do not configure the switch at all, the switch operates with these default settings:

- Default switch IP address, subnet mask, and default gateway is 0.0.0.0. For more information, see Chapter 4, "Performing Switch Setup Configuration," and Chapter 25, "Configuring DHCP."
- Default domain name is not configured. For more information, see Chapter 4, "Performing Switch Setup Configuration."
- DHCP client is enabled, the DHCP server is enabled (only if the device acting as a DHCP server is configured and is enabled), and the DHCP relay agent is enabled (only if the device is acting as a DHCP relay agent is configured and is enabled). For more information, see Chapter 4, "Performing Switch Setup Configuration," and Chapter 25, "Configuring DHCP."
- Switch cluster is disabled. For more information about switch clusters, see Chapter 6, "Configuring Switch Clusters," and the *Getting Started with Cisco Network Assistant*, available on Cisco.com.
- No passwords are defined. For more information, see Chapter 7, "Performing Switch Administration."
- System name and prompt is Switch. For more information, see Chapter 7, "Performing Switch Administration."
- NTP is enabled. For more information, see Chapter 7, "Performing Switch Administration."
- DNS is enabled. For more information, see Chapter 7, "Performing Switch Administration."
- TACACS+ is disabled. For more information, see Chapter 12, "Configuring Switch-Based Authentication."
- RADIUS is disabled. For more information, see Chapter 12, "Configuring Switch-Based Authentication."
- The standard HTTP server and Secure Socket Layer (SSL) HTTPS server are both enabled. For more information, see Chapter 12, "Configuring Switch-Based Authentication."
- IEEE 802.1x is disabled. For more information, see Chapter 13, "Configuring IEEE 802.1x Port-Based Authentication."
- Port parameters
  - Operating mode is Layer 2 (switch port). For more information, see Chapter 15, "Configuring Interface Characteristics."
  - Interface speed and duplex mode is autonegotiate. For more information, see Chapter 15, "Configuring Interface Characteristics."
  - Auto-MDIX is enabled. For more information, see Chapter 15, "Configuring Interface Characteristics."
  - Flow control is off. For more information, see Chapter 15, "Configuring Interface Characteristics."
- VLANs
  - Default VLAN is VLAN 1. For more information, see Chapter 17, "Configuring VLANs."
  - VLAN trunking setting is dynamic auto (DTP). For more information, see Chapter 17, "Configuring VLANs."
  - Trunk encapsulation is negotiate. For more information, see Chapter 17, "Configuring VLANs."
  - VTP mode is server. For more information, see Chapter 18, "Configuring VTP."
  - VTP version is Version 1. For more information, see Chapter 18, "Configuring VTP."
  - Voice VLAN is disabled. For more information, see Chapter 19, "Configuring Voice VLAN."
- STP, PVST+ is enabled on VLAN 1. For more information, see Chapter 20, "Configuring STP."

- MSTP is disabled. For more information, see Chapter 21, “Configuring MSTP.”
- Optional spanning-tree features are disabled. For more information, see Chapter 22, “Configuring Optional Spanning-Tree Features.”
- FlexLinks are not configured. For more information, see Chapter 24, “Configuring FlexLinks and the MAC Address-Table Move Update.”
- DHCP snooping is disabled. The DHCP snooping information option is enabled. For more information, see Chapter 25, “Configuring DHCP.”
- IP source guard is disabled. For more information, see Chapter 25, “Configuring DHCP.”
- DHCP server port-based address allocation is disabled. For more information, see Chapter 25, “Configuring DHCP.”
- Dynamic ARP inspection is disabled on all VLANs. For more information, see Chapter 26, “Configuring Dynamic ARP Inspection.”
- IGMP snooping is enabled. No IGMP filters are applied. For more information, see Chapter 28, “Configuring IGMP Snooping and MVR.”
- IGMP throttling setting is deny. For more information, see Chapter 28, “Configuring IGMP Snooping and MVR.”
- The IGMP snooping querier feature is disabled. For more information, see Chapter 28, “Configuring IGMP Snooping and MVR.”
- MVR is disabled. For more information, see Chapter 28, “Configuring IGMP Snooping and MVR.”
- Port-based traffic
  - Broadcast, multicast, and unicast storm control is disabled. For more information, see Chapter 29, “Configuring Port-Based Traffic Control.”
  - No protected ports are defined. For more information, see Chapter 29, “Configuring Port-Based Traffic Control.”
  - Unicast and multicast traffic flooding is not blocked. For more information, see Chapter 29, “Configuring Port-Based Traffic Control.”
  - No secure ports are configured. For more information, see Chapter 29, “Configuring Port-Based Traffic Control.”
- CDP is enabled. For more information, see Chapter 32, “Configuring CDP.”
- UDLD is disabled. For more information, see Chapter 33, “Configuring UDLD.”
- SPAN and RSPAN are disabled. For more information, see Chapter 30, “Configuring SPAN and RSPAN.”
- RMON is disabled. For more information, see Chapter 34, “Configuring RMON.”
- Syslog messages are enabled and appear on the console. For more information, see Chapter 35, “Configuring System Message Logging.”
- SNMP is enabled (Version 1). For more information, see Chapter 36, “Configuring SNMP.”
- No ACLs are configured. For more information, see Chapter 37, “Configuring Network Security with ACLs.”
- QoS is disabled. For more information, see Chapter 38, “Configuring Standard QoS.”
- No EtherChannels are configured. For more information, see Chapter 40, “Configuring EtherChannels.”
- IP unicast routing is disabled. For more information, see Chapter 41, “Configuring IP Unicast Routing.”

Network Configuration Examples

This section provides network configuration concepts and includes examples of using the switch to create dedicated network segments and interconnecting the segments through Fast Ethernet and Gigabit Ethernet connections.

- Design Concepts for Using the Switch, page 1-14
- Ethernet-to-the-Factory Architecture, page 1-15

Design Concepts for Using the Switch

As your network users compete for network bandwidth, it takes longer to send and receive data. When you configure your network, consider the bandwidth required by your network users and the relative priority of the network applications that they use.

Table 1-1 describes what can cause network performance to degrade and how you can configure your network to increase the bandwidth available to your network users.

Table 1-1 Increasing Network Performance

Network Demands	Suggested Design Methods
Too many users on a single network segment and a growing number of users accessing the Internet	• Create smaller network segments so that fewer users share the bandwidth, and use VLANs and IP subnets to place the network resources in the same logical network as the users who access those resources most. • Use full-duplex operation between the switch and its connected workstations.
• Increased power of new PCs, workstations, and servers • High bandwidth demand from networked applications (such as e-mail with large attached files) and from bandwidth-intensive applications (such as multimedia)	• Connect global resources, such as servers and routers to which the network users require equal access, directly to the high-speed switch ports so that they have their own high-speed segment. • Use the EtherChannel feature between the switch and its connected servers and routers.

Bandwidth alone is not the only consideration when designing your network. As your network traffic profiles evolve, consider providing network services that can support applications for voice and data integration, multimedia integration, application prioritization, and security. Table 1-2 describes some network demands and how you can meet them.

Table 1-2 Providing Network Services

Network Demands	Suggested Design Methods
Efficient bandwidth usage for multimedia applications and guaranteed bandwidth for critical applications	• Use IGMP snooping to efficiently forward multimedia and multicast traffic. • Use other QoS mechanisms such as packet classification, marking, scheduling, and congestion avoidance to classify traffic with the appropriate priority level, which provides maximum flexibility and support for mission-critical, unicast, and multicast and multimedia applications. • Use MVR to continuously send multicast streams in a multicast VLAN but to isolate the streams from subscriber VLANs for bandwidth and security reasons.
High demand on network redundancy and availability to provide *always on* mission-critical applications	• Use VLAN trunks and BackboneFast for traffic-load balancing on the uplink ports so that the uplink port with a lower relative port cost is selected to carry the VLAN traffic.
An evolving demand for IP telephony	• Use QoS to prioritize applications such as IP telephony during congestion and to help control both delay and jitter within the network. • Use switches that support at least two queues per port to prioritize voice and data traffic as either high- or low-priority, based on IEEE 802.1p/Q. The switch supports at least four queues per port. • Use voice VLAN IDs (VVIDs) to provide separate VLANs for voice traffic.

Ethernet-to-the-Factory Architecture

This section is an overview of the Ethernet-to-the-Factory (EttF) architecture that provides network and security services to the devices and applications in automation and control systems. It then integrates those into the wider enterprise network.

EttF architecture applies to many types of manufacturing environments, but it must be tailored to the industry type, the manufacturing type, and the production-facility size. Deployments can range from small networks (less than 50 devices), to medium-sized networks (less than 200 devices), and to large networks (up to and more than 1000 devices).

Within the EttF architecture are conceptual structures called *zones* that separate the various functions, from the highest-level enterprise switches and processes to the smallest devices that control more detailed processes and devices on the factory floor. See Figure 1-1.

For more information about EttF architecture, see this URL:

http://www.cisco.com/web/strategy/manufacturing/ettf_overview.html

Enterprise Zone

The *enterprise zone* comprises the centralized IT systems and functions. Wired and wireless access is available to enterprise network services, such as enterprise resource management, business-to-business, and business-to-customer services.The basic business administration tasks, such as site business planning and logistics, are performed here and rely on standard IT services. Guest access systems are often located here, although it is not uncommon to find them in lower levels of the framework to gain flexibility that might be difficult to achieve at the enterprise level.

Demilitarized Zone

The *demilitarized zone* (DMZ) provides a buffer for sharing of data and services between the enterprise and manufacturing zones. The DMZ maintains availability, addresses security vulnerabilities, and abiding by regulatory compliance mandates. The DMZ provides segmentation of organizational control, for example, between the IT and production organizations. Different policies for each organization can be applied and contained. For example, the production organization might apply security policies to the manufacturing zone that are different than those applied to the IT organization.

Manufacturing Zone

The *manufacturing zone* comprises the cell networks and site-level activities. All the systems, devices, and controllers that monitor the plant operations are in this zone. The cell zone is a functional area within a production facility.

The cell zone is a set of devices, controllers, and so on, that provide the real-time control of a functional aspect of the automation process. They are all in real-time communication with each other. This zone requires clear isolation and protection from the other levels of plant or enterprise operations.

Figure 1-1 shows the EttF architecture.

Figure 1-1 Ethernet-to-the-Factory Architecture

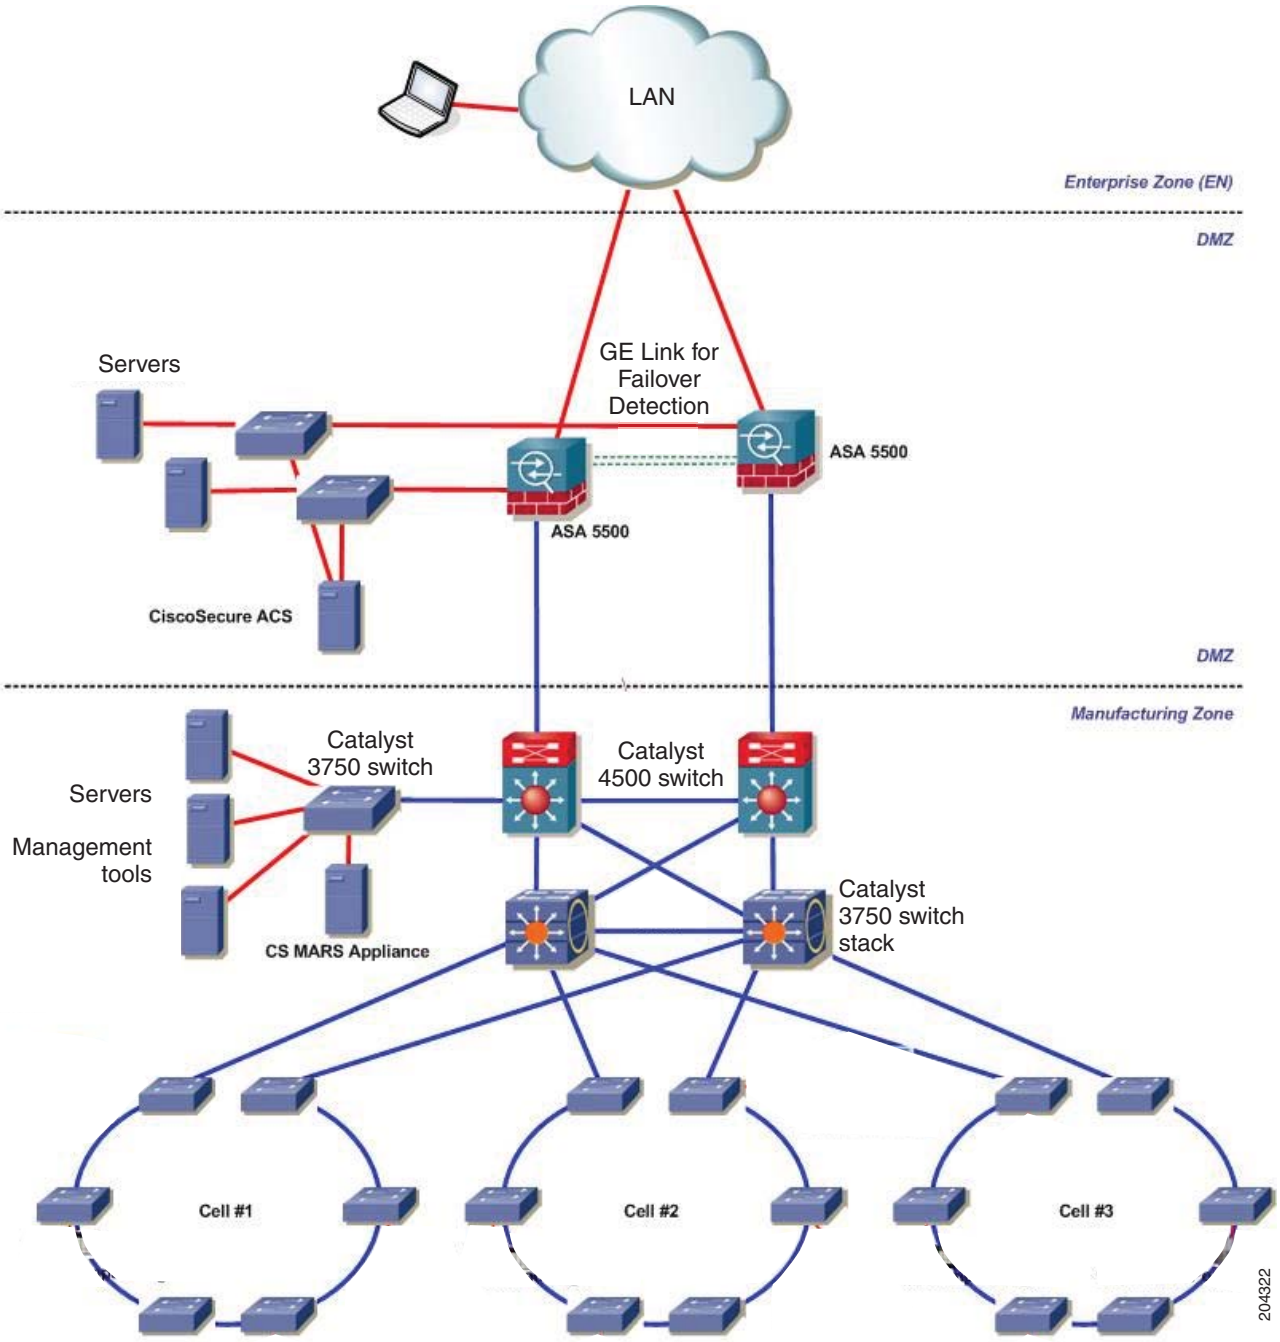

Topology Options

Topology design starts with considering how devices are connected to the network. The cell network also requires physical topologies that meet the physical constraints of the production floor. This section provides guidelines for topology designs and describes the trunk-drop, ring, and redundant-star topologies.

- Physical layout—The layout of the production environment drives the topology design. For example, a trunk-drop or ring topology is a good choice for a long conveyor-belt system, but a redundant-star configuration is not a good choice.
- Real-time communications—Latency and jitter are primarily caused by the amount of traffic and number of hops a packet must make to reach its destination. The amount of traffic in a Layer 2 network is driven by various factors, but the number of devices is important. Follow these guidelines for real-time communications:
  - The amount of latency introduced per Layer 2 hop should be considered. For instance, there is a higher latency with 100 Mb interfaces than there is with 1 Gigabit interfaces.
  - Bandwidth should not consistently exceed 50 percent of the interface capacity on any switch.
  - The CPU should not consistently exceed 50 to 70 percent utilization. Above this level, the switch might not properly process control packets and might behave abnormally.

These are the key connectivity considerations:

- Devices are connected to a switch through a single network connection or an IP-enabled I/O block or linking device if they do not support Ethernet. Most devices have no or limited failover capabilities and therefore cannot effectively use redundant network connections.
- Redundant connections can be used in certain industries and applications, such as process-related industries that are applied to critical infrastructure.

Cell Network—Trunk-Drop Topology

Switches are connected to each other to form a chain of switches in a *trunk-drop* topology (also known as a *cascaded* topology). See Figure 1-2.

- The connection between the Layer 3 switch and the first Layer 2 switch is very susceptible to oversubscription, which can degrade network performance.
- There is no redundancy to the loss of a connection.

Figure 1-2 Cell Network–Trunk-Drop Topology

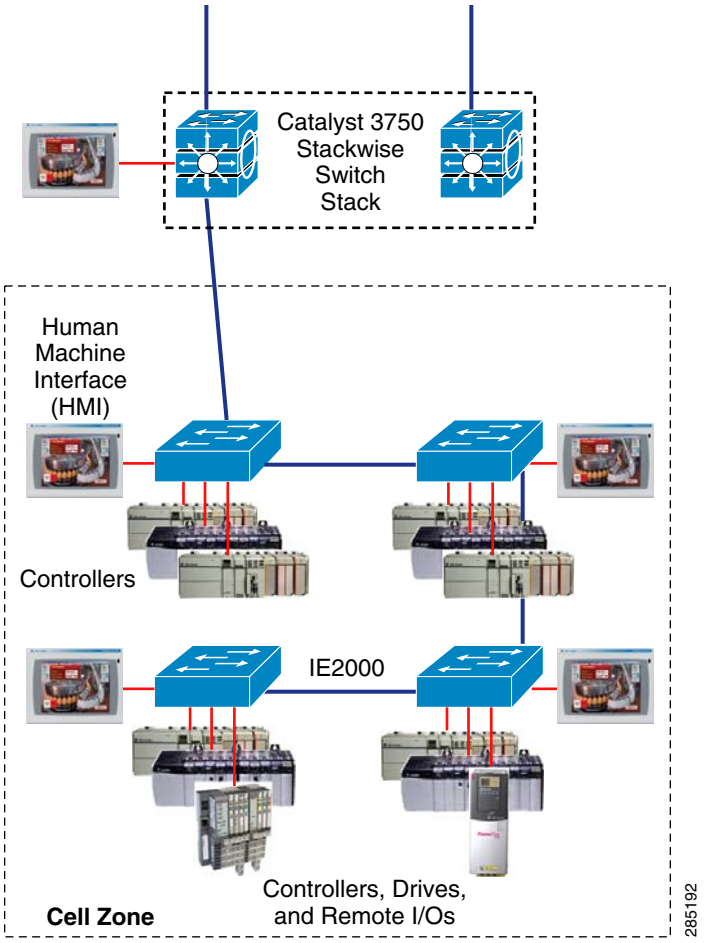

Cell Network—Ring Topology

A ring topology is similar to a trunk-drop topology except that the last switch in the chain is connected to the Layer 3 switch that forms a network ring. If a connection is lost in a ring, each switch maintains connectivity to the other switches. See Figure 1-3.

- The network can only recover from the loss of a single connection.
- It is more difficult to implement because it requires additional protocol implementation and Rapid Spanning Tree Protocol (RSTP).
- Although better than the trunk-drop, the top of the ring (connections to the Layer 3 switches) can become a bottleneck and is susceptible to oversubscription, which can degrade network performance.

Figure 1-3 Cell Network–Ring Topology

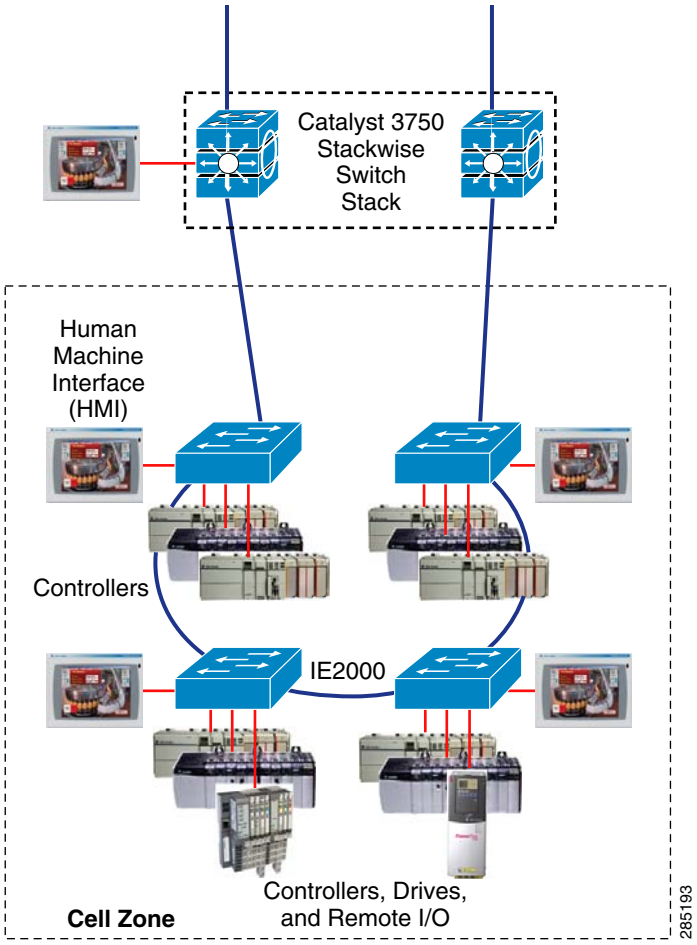

Cell Network—Redundant-Star Topology

In a redundant-star topology, every Layer 2 access switch has dual connections to a Layer 3 distribution switch. Devices are connected to the Layer 2 switches. See Figure 1-4.

- Any Layer 2 switch is always only two hops to another Layer 2 switch.
- In the Layer 2 network, each switch has dual connections to the Layer 3 devices.
- The Layer 2 network is maintained even if multiple connections are lost.

Figure 1-4 Cell Network–Redundant Star Topology

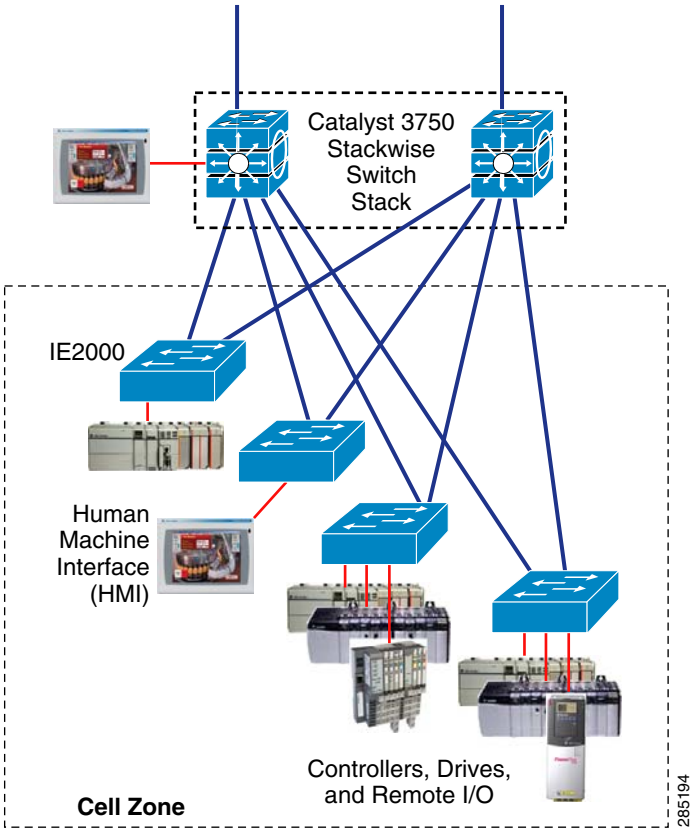

Where to Go Next

Before configuring the switch, review these sections for startup information:

- Chapter 2, "Using the Command-Line Interface"
- Chapter 4, "Performing Switch Setup Configuration"

To locate and download MIBs for a specific Cisco product and release, use the Cisco MIB Locator: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml.

CHAPTER 2

Using the Command-Line Interface

Information About Using the Command-Line Interface

This chapter describes the Cisco IOS command-line interface (CLI) and how to use it to configure your switch.

Command Modes

The Cisco IOS user interface is divided into many different modes. The commands available to you depend on which mode you are currently in. Enter a question mark (?) at the system prompt to obtain a list of commands available for each command mode.

When you start a session on the switch, you begin in user mode, often called user EXEC mode. Only a limited subset of the commands are available in user EXEC mode. For example, most of the user EXEC commands are one-time commands, such as **show** commands, which show the current configuration status, and **clear** commands, which clear counters or interfaces. The user EXEC commands are not saved when the switch reboots.

To have access to all commands, you must enter privileged EXEC mode. You must enter a password to enter privileged EXEC mode. From this mode, you can enter any privileged EXEC command or enter global configuration mode.

Using the configuration modes (global, interface, and line), you can make changes to the running configuration. If you save the configuration, these commands are stored and used when the switch reboots. To access the various configuration modes, you must start at global configuration mode. From global configuration mode, you can enter interface configuration mode and line configuration mode.

Table 2-1 describes the main command modes, how to access each one, the prompt you see in that mode, and how to exit the mode. The examples in the table use the hostname Switch.

Table 2-1 Command Mode Summary

Mode	Access Method	Prompt	Exit Method	About This Mode
User EXEC	Begin a session with your switch.	`Switch>`	Enter **logout** or **quit**.	Use this mode to • Change terminal settings. • Perform basic tests. • Display system information.
Privileged EXEC	While in user EXEC mode, enter the **enable** command.	`Switch#`	Enter **disable** to exit.	Use this mode to verify commands that you have entered. Use a password to protect access to this mode.
Global configuration	While in privileged EXEC mode, enter the **configure** command.	`Switch(config)#`	To exit to privileged EXEC mode, enter **exit** or **end**, or press **Ctrl-Z**.	Use this mode to configure parameters that apply to the entire switch.
Config-vlan	While in global configuration mode, enter the **vlan** *vlan-id* command.	`Switch(config-vlan)#`	To exit to global configuration mode, enter the **exit** command. To return to privileged EXEC mode, press **Ctrl-Z** or enter **end**.	Use this mode to configure VLAN parameters. When VTP mode is transparent, you can create extended-range VLANs (VLAN IDs greater than 1005) and save configurations in the switch startup configuration file.
VLAN configuration	While in privileged EXEC mode, enter the **vlan database** command.	`Switch(vlan)#`	To exit to privileged EXEC mode, enter **exit**.	Use this mode to configure VLAN parameters for VLANs 1 to 1005 in the VLAN database.

Table 2-1 Command Mode Summary (continued)

Mode	Access Method	Prompt	Exit Method	About This Mode
Interface configuration	While in global configuration mode, enter the **interface** command (with a specific interface).	`Switch(config-if)#`	To exit to global configuration mode, enter **exit**. To return to privileged EXEC mode, press **Ctrl-Z** or enter **end**.	Use this mode to configure parameters for the Ethernet ports. For information about defining interfaces, see the "Using Interface Configuration Mode" section on page 15-6. To configure multiple interfaces with the same parameters, see the "Configuring a Range of Interfaces" section on page 15-13.
Line configuration	While in global configuration mode, specify a line with the **line vty** or **line console** command.	`Switch(config-line)#`	To exit to global configuration mode, enter **exit**. To return to privileged EXEC mode, press **Ctrl-Z** or enter **end**.	Use this mode to configure parameters for the terminal line.

For more detailed information on the command modes, see the command reference guide for this release.

Help System

You can enter a question mark (?) at the system prompt to display a list of commands available for each command mode. You can also obtain a list of associated keywords and arguments for any command, as shown in Table 2-2.

Table 2-2 Help Summary

Command	Purpose
help	Obtain a brief description of the help system in any command mode.

*abbreviated-command-entry***?**	Obtain a list of commands that begin with a particular character string.
	For example:
	`Switch# di?`
	`dir disable disconnect`

abbreviated-command-entry<**Tab**>	Complete a partial command name.
	For example:
	`Switch# sh conf<tab>`
	`Switch# show configuration`

Table 2-2 Help Summary (continued)

Command	Purpose
?	List all commands available for a particular command mode. For example: `Switch> ?`
command **?**	List the associated keywords for a command. For example: `Switch> show ?`

command keyword **?**	List the associated arguments for a keyword.
	For example:
	`Switch(config)# cdp holdtime ?`
	`  <10-255> Length of time (in sec) that receiver must keep this packet`

Understanding Abbreviated Commands

You need to enter only enough characters for the switch to recognize the command as unique.

This example shows how to enter the **show configuration** privileged EXEC command in an abbreviated form:

```
Switch# show conf
```

No and default Forms of Commands

Almost every configuration command also has a **no** form. In general, use the **no** form to disable a feature or function or reverse the action of a command. For example, the **no shutdown** interface configuration command reverses the shutdown of an interface. Use the command without the keyword **no** to reenable a disabled feature or to enable a feature that is disabled by default.

Configuration commands can also have a **default** form. The **default** form of a command returns the command setting to its default. Most commands are disabled by default, so the **default** form is the same as the **no** form. However, some commands are enabled by default and have variables set to certain default values. In these cases, the **default** command enables the command and sets variables to their default values.

CLI Error Messages

Table 2-3 lists some error messages that you might encounter while using the CLI to configure your switch.

Table 2-3 Common CLI Error Messages

Error Message	Meaning	How to Get Help
`% Ambiguous command: "show con"`	You did not enter enough characters for your switch to recognize the command.	Reenter the command followed by a question mark (?) with a space between the command and the question mark. The possible keywords that you can enter with the command appear.
`% Incomplete command.`	You did not enter all the keywords or values required by this command.	Reenter the command followed by a question mark (?) with a space between the command and the question mark. The possible keywords that you can enter with the command appear.
`% Invalid input detected at '^' marker.`	You entered the command incorrectly. The caret (^) marks the point of the error.	Enter a question mark (?) to display all the commands that are available in this command mode. The possible keywords that you can enter with the command appear.

Configuration Logging

You can log and view changes to the switch configuration. You can use the Configuration Change Logging and Notification feature to track changes on a per-session and per-user basis. The logger tracks each configuration command that is applied, the user who entered the command, the time that the command was entered, and the parser return code for the command. This feature includes a mechanism for asynchronous notification to registered applications whenever the configuration changes. You can choose to have the notifications sent to the syslog.

Note Only CLI or HTTP changes are logged.

How to Use the CLI to Configure Features

Configuring the Command History

The software provides a history or record of commands that you have entered. The command history feature is particularly useful for recalling long or complex commands or entries, including access lists. You can customize this feature to suit your needs as described in these sections:

- Changing the Command History Buffer Size, page 2-6 (optional)
- Recalling Commands, page 2-6 (optional)
- Disabling the Command History Feature, page 2-7 (optional)

Changing the Command History Buffer Size

By default, the switch records ten command lines in its history buffer. You can alter this number for a current terminal session or for all sessions on a particular line. These procedures are optional.

Beginning in privileged EXEC mode, enter this command to change the number of command lines that the switch records during the current terminal session:

```
Switch# terminal history [size number-of-lines]
```

The range is from 0 to 256.

Beginning in line configuration mode, enter this command to configure the number of command lines the switch records for all sessions on a particular line:

```
Switch(config-line)# history [size number-of-lines]
```

The range is from 0 to 256.

Recalling Commands

To recall commands from the history buffer, perform one of the actions listed in Table 2-4. These actions are optional.

Table 2-4 Recalling Commands

Action[1]	Result
Press **Ctrl-P** or the up arrow key.	Recall commands in the history buffer, beginning with the most recent command. Repeat the key sequence to recall successively older commands.
Press **Ctrl-N** or the down arrow key.	Return to more recent commands in the history buffer after recalling commands with **Ctrl-P** or the up arrow key. Repeat the key sequence to recall successively more recent commands.
show history	While in privileged EXEC mode, list the last several commands that you just entered. The number of commands that appear is controlled by the setting of the **terminal history** global configuration command and the **history** line configuration command.

1. The arrow keys function only on ANSI-compatible terminals such as VT100s.

Disabling the Command History Feature

The command history feature is automatically enabled. You can disable it for the current terminal session or for the command line. These procedures are optional.

To disable the feature during the current terminal session, enter the **terminal no history** privileged EXEC command.

To disable command history for the line, enter the **no history** line configuration command.

Using Editing Features

This section describes the editing features that can help you manipulate the command line. It contains these sections:

- Enabling and Disabling Editing Features, page 2-7 (optional)
- Editing Commands Through Keystrokes, page 2-7 (optional)
- Editing Command Lines That Wrap, page 2-9 (optional)

Enabling and Disabling Editing Features

Although enhanced editing mode is automatically enabled, you can disable it, reenable it, or configure a specific line to have enhanced editing. These procedures are optional.

To globally disable enhanced editing mode, enter this command in line configuration mode:

```
Switch (config-line)# no editing
```

To reenable the enhanced editing mode for the current terminal session, enter this command in privileged EXEC mode:

```
Switch# terminal editing
```

To reconfigure a specific line to have enhanced editing mode, enter this command in line configuration mode:

```
Switch(config-line)# editing
```

Editing Commands Through Keystrokes

Table 2-5 shows the keystrokes that you need to edit command lines. These keystrokes are optional.

Table 2-5 Editing Commands through Keystrokes

Capability	Keystroke[1]	Purpose
Move around the command line to make changes or corrections.	Press **Ctrl-B**, or press the left arrow key.	Move the cursor back one character.

Table 2-5 Editing Commands through Keystrokes (continued)

Capability	Keystroke[1]	Purpose
	Press **Ctrl-F**, or press the right arrow key.	Move the cursor forward one character.
	Press **Ctrl-A**.	Move the cursor to the beginning of the command line.
	Press **Ctrl-E**.	Move the cursor to the end of the command line.
	Press **Esc B**.	Move the cursor back one word.
	Press **Esc F**.	Move the cursor forward one word.
	Press **Ctrl-T**.	Transpose the character to the left of the cursor with the character located at the cursor.
Recall commands from the buffer and paste them in the command line. The switch provides a buffer with the last ten items that you deleted.	Press **Ctrl-Y.**	Recall the most recent entry in the buffer.
	Press **Esc Y**.	Recall the next buffer entry. The buffer contains only the last 10 items that you have deleted or cut. If you press **Esc Y** more than ten times, you cycle to the first buffer entry.
Delete entries if you make a mistake or change your mind.	Press the **Delete** or **Backspace** key.	Erase the character to the left of the cursor.
	Press **Ctrl-D**.	Delete the character at the cursor.
	Press **Ctrl-K**.	Delete all characters from the cursor to the end of the command line.
	Press **Ctrl-U** or **Ctrl-X**.	Delete all characters from the cursor to the beginning of the command line.
	Press **Ctrl-W**.	Delete the word to the left of the cursor.
	Press **Esc D**.	Delete from the cursor to the end of the word.
Capitalize or lowercase words or capitalize a set of letters.	Press **Esc C**.	Capitalize at the cursor.
	Press **Esc L**.	Change the word at the cursor to lowercase.
	Press **Esc U**.	Capitalize letters from the cursor to the end of the word.
Designate a particular keystroke as an executable command, perhaps as a shortcut.	Press **Ctrl-V** or **Esc Q**.	

Table 2-5 Editing Commands through Keystrokes (continued)

Capability	Keystroke[1]	Purpose
Scroll down a line or screen on displays that are longer than the terminal screen can display. **Note** The More prompt is used for any output that has more lines than can be displayed on the terminal screen, including **show** command output. You can use the **Return** and **Space** bar keystrokes whenever you see the More prompt.	Press the **Return** key.	Scroll down one line.
	Press the **Space** bar.	Scroll down one screen.
Redisplay the current command line if the switch suddenly sends a message to your screen.	Press **Ctrl-L** or **Ctrl-R.**	Redisplay the current command line.

1. The arrow keys function only on ANSI-compatible terminals such as VT100s.

Editing Command Lines That Wrap

You can use a wraparound feature for commands that extend beyond a single line on the screen. When the cursor reaches the right margin, the command line shifts ten spaces to the left. You cannot see the first ten characters of the line, but you can scroll back and check the syntax at the beginning of the command. The keystroke actions are optional.

To scroll back to the beginning of the command entry, press **Ctrl-B** or the left arrow key repeatedly. You can also press **Ctrl-A** to immediately move to the beginning of the line.

The arrow keys function only on ANSI-compatible terminals such as VT100s.

In this example, the **access-list** global configuration command entry extends beyond one line. When the cursor first reaches the end of the line, the line is shifted ten spaces to the left and redisplayed. The dollar sign ($) shows that the line has been scrolled to the left. Each time the cursor reaches the end of the line, the line is again shifted ten spaces to the left.

```
Switch(config)# access-list 101 permit tcp 131.108.2.5 255.255.255.0 131.108.1
Switch(config)# $ 101 permit tcp 131.108.2.5 255.255.255.0 131.108.1.20 255.25
Switch(config)# $t tcp 131.108.2.5 255.255.255.0 131.108.1.20 255.255.255.0 eq
Switch(config)# $108.2.5 255.255.255.0 131.108.1.20 255.255.255.0 eq 45
```

After you complete the entry, press **Ctrl-A** to check the complete syntax before pressing the **Return** key to execute the command. The dollar sign ($) appears at the end of the line to show that the line has been scrolled to the right:

```
Switch(config)# access-list 101 permit tcp 131.108.2.5 255.255.255.0 131.108.1$
```

The software assumes you have a terminal screen that is 80 columns wide. If you have a different width, use the **terminal width** privileged EXEC command to set the width of your terminal.

Use line wrapping with the command history feature to recall and modify previous complex command entries. For information about recalling previous command entries, see the "Editing Commands Through Keystrokes" section on page 2-7.

Searching and Filtering Output of show and more Commands

You can search and filter the output for **show** and **more** commands. This is useful when you need to sort through large amounts of output or if you want to exclude output that you do not need to see. Using these commands is optional.

To use this functionality, enter a **show** or **more** command followed by the pipe character (|), one of the keywords **begin**, **include**, or **exclude**, and an expression that you want to search for or filter out:

command **|** {**begin** | **include** | **exclude**} *regular-expression*

Expressions are case sensitive. For example, if you enter **| exclude output**, the lines that contain **output** are not displayed, but the lines that contain **Output** appear.

This example shows how to include in the output display only lines where the expression **protocol** appears:

```
Switch# show interfaces | include protocol
Vlan1 is up, line protocol is up
Vlan10 is up, line protocol is down
```

Accessing the CLI

You can access the CLI through a console connection, through Telnet, or by using the browser.

Accessing the CLI through a Console Connection or through Telnet

Before you can access the CLI, you must connect a terminal or PC to the switch console port and power on the switch, as described in the getting started guide that shipped with your switch. Then, to understand the boot process and the options available for assigning IP information, see Chapter 4, "Performing Switch Setup Configuration."

If your switch is already configured, you can access the CLI through a local console connection or through a remote Telnet session, but your switch must first be configured for this type of access. For more information, see the "Setting a Telnet Password for a Terminal Line" section on page 12-28.

You can use one of these methods to establish a connection with the switch:

- Connect the switch console port to a management station or dial-up modem. For information about connecting to the console port, see the switch getting started guide or hardware installation guide.
- Use any Telnet TCP/IP or encrypted Secure Shell (SSH) package from a remote management station. The switch must have network connectivity with the Telnet or SSH client, and the switch must have an enable secret password configured.

  For information about configuring the switch for Telnet access, see the "Setting a Telnet Password for a Terminal Line" section on page 12-28. The switch supports up to 16 simultaneous Telnet sessions. Changes made by one Telnet user are reflected in all other Telnet sessions.

  For information about configuring the switch for SSH, see the "Configuring the SSH Server" section on page 12-40. The switch supports up to five simultaneous secure SSH sessions.

After you connect through the console port, through a Telnet session or through an SSH session, the user EXEC prompt appears on the management station.

CHAPTER 3

Configuring Switch Alarms

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Information About Switch Alarms

The switch software monitors switch conditions on a per-port or a switch basis. If the conditions present on the switch or a port do not match the set parameters, the switch software triggers an alarm or a system message. By default, the switch software sends the system messages to a system message logging facility, or a *syslog* facility. You can also configure the switch to send Simple Network Management Protocol (SNMP) traps to an SNMP server. You can configure the switch to trigger an external alarm device by using the alarm relay.

Global Status Monitoring Alarms

The switch processes alarms related to temperature and power supply conditions, referred to as global or facility alarms.

Table 3-1 Global Status Monitoring Alarms

Alarm	Description
Power supply alarm	By default, the switch monitors a single power supply. If you configure a dual power supply, an alarm triggers if one power supply fails. You can configure the power supply alarm to be connected to the hardware relays. For more information, see the "Configuring the Power Supply Alarms" section on page 3-6.
Temperature alarms	The switch contains one temperature sensor with a primary and secondary temperature setting. The sensor monitors the environmental conditions inside the switch. The primary and secondary temperature alarms can be set as follows: • The primary alarm is enabled automatically to trigger both at a low temperature, –4°F (–20°C) and a high temperature, 203°F (95°C). It cannot be disabled. By default, the primary temperature alarm is associated with the major relay. • The secondary alarm triggers when the system temperature is higher or lower than the configured high and low temperature thresholds. The secondary alarm is disabled by default. For more information, see the "Configuring the Switch Temperature Alarms" section on page 3-6.
SD-Card	By default the alarm is disabled.

FCS Error Hysteresis Threshold

The Ethernet standard calls for a maximum bit-error rate of 10^{-8}. The bit error-rate range is from 10^{-6} to 10^{-11}. The bit error-rate input to the switch is a positive exponent. If you want to configure the bit error-rate of 10^{-9}, enter the value 9 for the exponent. By default, the FCS bit error-rate is 10^{-8}.

You can set the FCS error hysteresis threshold to prevent the toggle of the alarm when the actual bit-error rate fluctuates near the configured rate. The hysteresis threshold is defined as the ratio between the alarm clear threshold to the alarm set threshold, expressed as a percentage value.

For example, if the FCS bit error-rate alarm value is configured to 10^{-8}, that value is the alarm set threshold. To set the alarm clear threshold at $5*10^{-10}$, the hysteresis, value *h*, is determined as follows:

h = alarm clear threshold / alarm set threshold

$h = 5*10^{-10} / 10^{-8} = 5*10^{-2} = 0.05 = 5$ percent

The FCS hysteresis threshold is applied to all ports on the switch. The allowable range is from 1 to 10 percent. The default value is 10 percent. See the "Configuring the FCS Bit Error Rate Alarm" section on page 3-7 for more information.

Port Status Monitoring Alarms

The switch can also monitor the status of the Ethernet ports and generate alarm messages based on the alarms listed in Table 3-2. To save user time and effort, it supports changeable alarm configurations by using alarm profiles. You can create a number of profiles and assign one of these profiles to each Ethernet port.

Alarm profiles provide a mechanism for you to enable or disable alarm conditions for a port and associate the alarm conditions with one or both alarm relays. You can also use alarm profiles to set alarm conditions to send alarm traps to an SNMP server and system messages to a syslog server. The alarm profile *defaultPort* is applied to all interfaces in the factory configuration (by default).

Note You can associate multiple alarms to one relay or one alarm to both relays.

Table 3-2 lists the port status monitoring alarms and their descriptions and functions. Each fault condition is assigned a severity level based on the Cisco IOS System Error Message Severity Level.

Table 3-2 Port Status Monitoring Alarms

Alarm List ID	Alarm	Description
1	Link Fault alarm	The switch generates a link fault alarm when problems with a port physical layer cause unreliable data transmission. A typical link fault condition is loss of signal or clock. The link fault alarm is cleared automatically when the link fault condition is cleared. The severity for this alarm is *error condition*, level 3.
2	Port not Forwarding alarm	The switch generates a port not-forwarding alarm when a port is not forwarding packets. This alarm is cleared automatically when the port begins to forward packets. The severity for this alarm is *warning*, level 4.
3	Port not Operating alarm	The switch generates a port not-operating alarm when a port fails during the startup self-test. When triggered, the port not-operating alarm is only cleared when the switch is restarted and the port is operational. The severity for this alarm is *error condition*, level 3.
4	FCS Bit Error Rate alarm	The switch generates an FCS bit error-rate alarm when the actual FCS bit error-rate is close to the configured rate. You can set the FCS bit error-rate by using the interface configuration CLI for each of the ports. See the "Configuring the FCS Bit Error Rate Alarm" section on page 3-7 for more information. The severity for this alarm is *error condition*, level 3.

Triggering Alarm Options

The switch supports these methods for triggering alarms:

- Configurable Relay

  The switch is equipped with one independent alarm relay that can be triggered by alarms for global, port status and SD flash card conditions. You can configure the relay to send a fault signal to an external alarm device, such as a bell, light, or other signaling device. You can associate any alarm condition with the alarm relay. Each fault condition is assigned a severity level based on the Cisco IOS System Error Message Severity Level.

  See the "Configuring the Power Supply Alarms" section on page 3-6 for more information on configuring the relay.

- SNMP Traps

  SNMP is an application-layer protocol that provides a message format for communication between managers and agents. The SNMP system consists of an SNMP manager, an SNMP agent, and a management information base (MIB).

The **snmp-server enable traps** command can be changed so that the user can send alarm traps to an SNMP server. You can use alarm profiles to set environmental or port status alarm conditions to send SNMP alarm traps. See the "Enabling SNMP Traps" section on page 3-9 for more information.

- Syslog Messages

  You can use alarm profiles to send system messages to a syslog server. See the "Configuring the Power Supply Alarms" section on page 3-6 for more information.

External Alarms

The switch supports two alarm inputs and one alarm output. The alarm input circuit is designed to sense if a dry contact is open or closed relative to the Alarm-In reference pin. The Alarm_Out is a relay with Normally Open and Normally Closed contacts. The switch software is configured to detect faults which are used to energize the relay coil and change the state on both of the relay contacts. Normally open contacts close and normally closed contacts open.

- **Open** means that the normal condition has current flowing through the contact (normally closed contact). The alarm is generated when the current stops flowing.
- **Closed** means that no current flows through the contact (normally open contact). The alarm is generated when current does flow.

Note Software can program the Alarm_In to trigger an alarm with either Open or Closed setting.

The alarm connector is a 6-pin screw terminal. This table lists pinouts for the alarm ports.

Pin #	Signal Name	Description
6	Alarm_Out_NO	Alarm output relay normally open contact
5	Alarm_Out_Com	Alarm output relay common contact
4	Alarm_Out-NC	Alarm output relay normally closed contact
3	Alarm_In2	Alarm input #2
2	Alarm_In_Ref	Alarm input reference
1	Alarm_In1	Alarm input #1

You can set the alarm severity to major, minor, or none. The severity is included in the alarm message and also sets the LED color when the alarm is triggered. The LED is red for a minor alarm and blinking red for a major alarm. If not set, the default alarm severity is minor.

For detailed information about the alarm connector, LEDs, alarm circuit and wiring installation, alarm ratings and ports, see the *Cisco IE 2000 Switch Hardware Installation Guide*.

Default Switch Alarm Settings

Table 3-3 Default Switch Alarm Settings

	Alarm	Default Setting
Global	Power supply alarm	Enabled in switch single power mode. No alarm. In dual-power supply mode, the default alarm notification is a system message to the console.
	Primary temperature alarm	Enabled for switch temperature range of 203°F (95°C) maximum to –4°F (–20°C) minimum. The primary switch temperature alarm is associated with the major relay.
	Secondary temperature alarm	Disabled.
	Output relay mode alarm	Normally deenergized. The alarm output has switched off or is in an off state.
Port	Link fault alarm	Disabled on all interfaces.
	Port not forwarding alarm	Disabled on all interfaces.
	Port not operating alarm	Enabled on all interfaces.
	FCS bit error rate alarm	Disabled on all interfaces.

How to Configure Switch Alarms

Configuring External Alarms

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**alarm contact** *contact-number* **description** *string*	(Optional) Configures a description for the alarm contact number. • The *contact-number* value is from 1 to 4. • The description string is up to 80 alphanumeric characters in length and is included in any generated system messages.
Step 3	**alarm contact** {*contact-number* | **all**} {**severity** { **major** | **minor** | **none**} | **trigger** {**closed** | **open**}}	Configures the trigger and severity for an alarm contact number or for all contact numbers. • Enter a contact number (1 to 4) or specify that you are configuring **all** alarms. • For **severity**, enter **major**, **minor** or **none**. If you do not configure a severity, the default is **minor**. • For **trigger**, enter **open** or **closed**. If you do not configure a trigger, the alarm is triggered when the circuit is **closed**.
Step 4	**alarm relay-mode energized**	(Optional) Configures the output relay mode to energized.
Step 5	**end**	Returns to privileged EXEC mode.

	Command	Purpose
Step 6	**show env alarm-contact**	Shows the configured alarm contacts.
Step 7	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring the Power Supply Alarms

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**power-supply dual**	Configures dual power supplies.
Step 3	**alarm facility power-supply disable**	Disables the power supply alarm.
Step 4	**alarm facility power-supply relay major**	Associates the power supply alarm to the relay.
Step 5	**alarm facility power-supply notifies**	Sends power supply alarm traps to an SNMP server.
Step 6	**alarm facility power-supply syslog**	Sends power supply alarm traps to a syslog server.
Step 7	**end**	Returns to privileged EXEC mode.
Step 8	**show env power**	Displays the switch power status.
Step 9	**show facility-alarm status**	Displays all generated alarms for the switch.
Step 10	**show alarm settings**	Verifies the configuration.
Step 11	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring the Switch Temperature Alarms

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**alarm facility temperature** {**primary** | **secondary**} **high** *threshold*	Sets the high temperature threshold value. Set the threshold from –238°F (–150°C) to 572°F (300°C).
Step 3	**alarm facility temperature primary low** *threshold*	Sets the low temperature threshold value. Set the threshold from –328°F (–200°C) to 482°F (250°C).
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show alarm settings**	Verifies the configuration.
Step 6	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Associating the Temperature Alarms to a Relay

By default, the primary temperature alarm is associated to the relay. You can use the **alarm facility temperature** global configuration command to associate the primary temperature alarm to an SNMP trap, or a syslog message, or to associate the secondary temperature alarm to the relay, an SNMP trap, or a syslog message.

Note The single relay on the switch is called the major relay.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**alarm facility temperature {primary | secondary} relay major**	Associates the primary or secondary temperature alarm to the relay.
Step 3	**alarm facility temperature {primary | secondary} notifies**	Sends primary or secondary temperature alarm traps to an SNMP server.
Step 4	**alarm facility temperature {primary | secondary} syslog**	Sends primary or secondary temperature alarm traps to a syslog server. Uses the **no alarm facility temperature secondary** command to disable the secondary temperature alarm.
Step 5	**end**	Returns to privileged EXEC mode.
Step 6	**show alarm settings**	Verifies the configuration.
Step 7	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring the FCS Bit Error Rate Alarm

Setting the FCS Error Threshold

The switch generates an FCS bit error-rate alarm when the actual rate is close to the configured rate.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Enters the interface to be configured, and enters interface configuration mode.
Step 3	**fcs-threshold** *value*	Sets the FCS error rate. For *value*, the range is 6 to 11 to set a maximum bit error rate of 10^{-6} to 10^{-11}. By default, the FCS bit error rate is 10^{-8}.
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show fcs-threshold**	Verifies the setting.
Step 6	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Setting the FCS Error Hysteresis Threshold

The hysteresis setting prevents the toggle of an alarm when the actual bit error-rate fluctuates near the configured rate. The FCS hysteresis threshold is applied to all ports of a switch.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**alarm facility fcs-hysteresis** *percentage*	Sets the hysteresis percentage for the switch. For *percentage*, the range is 1 to 10. The default value is 10 percent.
Step 3	**end**	Returns to privileged EXEC mode.
Step 4	**show running config**	Verifies the configuration.
Step 5	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring Alarm Profiles

Creating an Alarm Profile

You can use the **alarm profile** global configuration command to create an alarm profile or to modify an existing profile. When you create a new alarm profile, none of the alarms are enabled.

Note The only alarm enabled in the *defaultPort* profile is the Port not operating alarm.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**alarm profile** *name*	Creates the new profile or identifies an existing profile, and enters alarm profile configuration mode.
Step 3	**end**	Returns to privileged EXEC mode.
Step 4	**show alarm profile** *name*	Verifies the configuration.
Step 5	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Modifying an Alarm Profile

You can modify an alarm profile from alarm profile configuration mode.

You can enter more than one alarm type separated by a space.

Command	Purpose
alarm {fcs-error | link-fault | not-forwarding | not-operating}	(Optional) Adds or modifies alarm parameters for a specific alarm.
notifies {fcs-error | link-fault | not-forwarding | not-operating}	(Optional) Configures the alarm to send an SNMP trap to an SNMP server.

Command	Purpose
relay-major {fcs-error | link-fault | not-forwarding | not-operating}	(Optional) Configures the alarm to send an alarm trap to the relay.
syslog {fcs-error | link-fault | not-forwarding | not-operating}	(Optional) Configures the alarm to send an alarm trap to a syslog server.

Attaching an Alarm Profile to a Specific Port

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *port interface*	Enters interface configuration mode.
Step 3	**alarm-profile** ***name***	Attaches the specified profile to the interface.
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show alarm profile**	Verifies the configuration.
Step 6	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Enabling SNMP Traps

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**snmp-server enable traps alarms**	Enables the switch to send SNMP traps.
Step 3	**end**	Returns to privileged EXEC mode.
Step 4	**show alarm settings**	Verifies the configuration.
Step 5	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Monitoring and Maintaining Switch Alarms Status

Table 3-4 Commands for Displaying Global and Port Alarm Status

Command	Purpose
show alarm description ports	Displays an alarm number and its text description.
show alarm profile [*name*]	Displays all alarm profiles in the system or a specified profile.
show alarm settings	Displays all global alarm settings on the switch.
show env {**alarm-contact** | **all** | **power** | **temperature**}	Displays the status of environmental facilities on the switch.
show facility-alarm status [**critical** | **info** | **major** | **minor**]	Displays generated alarms on the switch.

Configuration Examples for Switch Alarms

Configuring External Alarms: Example

This example configures alarm input 1 named *door sensor* to assert a major alarm when the door circuit is closed and then displays the status and configuration for all alarms:

```
Switch(config)# alarm contact 1 description door sensor
Switch(config)# alarm contact 1 severity major
Switch(config)# alarm contact 1 trigger closed
Switch(config)# end
Switch(config)# show env alarm-contact
Switch# show env alarm-contact

ALARM CONTACT 1
   Status:      not asserted
   Description: door sensor
   Severity:    major
   Trigger:     closed
ALARM CONTACT 2
   Status:      not asserted
   Description: external alarm contact 2
   Severity:    minor
   Trigger:     closed
```

Associating Temperature Alarms to a Relay: Examples

This example sets the secondary temperature alarm to the major relay, with a high temperature threshold value of 113°F (45°C). All alarms and traps associated with this alarm are sent to a syslog server and an SNMP server.

```
Switch(config) # alarm facility temperature secondary high 45
Switch(config) # alarm facility temperature secondary relay major
Switch(config) # alarm facility temperature secondary syslog
Switch(config) # alarm facility temperature secondary notifies
```

This example sets the first (primary) temperature alarm to the major relay. All alarms and traps associated with this alarm are sent to a syslog server.

```
Switch(config) # alarm facility temperature primary syslog
Switch(config) # alarm facility temperature primary relay major
```

Creating or Modifying an Alarm Profile: Example

This example creates or modifies the alarm profile *fastE* for the Fast Ethernet port with link-down (*alarmList* ID 3) alarm enabled. The link-down alarm is connected to the major relay. This alarm also send notifications to an SNMP server and sends system messages to a syslog server.

```
Switch(config)# alarm profile fastE
Switch(config-alarm-profile)# alarm fcs-error
Switch(config-alarm-profile)# relay major link-fault
Switch(config-alarm-profile)# notifies not-forwarding
Switch(config-alarm-profile)# syslog not-forwarding
```

Setting the FCS Error Hysteresis Threshold: Example

This example shows how to set the FCS bit error rate for a port to 10^{-10}:

```
Switch# configure terminal
Switch(config)# interface fastethernet1/1
Switch(config-if) # fcs-threshold 10
```

Configuring a Dual Power Supply: Examples

This example shows how to configure two power supplies:

```
Switch# configure terminal
Switch(config)# power-supply dual
```

These examples show how to display information when two power supplies are not present which results in a triggered alarm.

```
Switch# show facility-alarm status
Source Severity Description Relay Time
Switch MAJOR 5 Redundant Pwr missing or failed NONE Mar 01
1993 00:23:52

Switch# show env power
POWER SUPPLY A is DC OK
POWER SUPPLY B is DC FAULTY <--

Switch# show hard led
SWITCH: 1
SYSTEM: GREEN
ALARM : ALT_RED_BLACK <--
```

Displaying Alarm Settings: Example

```
Switch# show alarm settings
Alarm relay mode: De-energized
Power Supply
         Alarm                   Enabled
         Relay
         Notifies                Disabled
         Syslog                  Enabled
Temperature-Primary
         Alarm                   Enabled
         Thresholds              MAX: 95C                MIN:  -20C
         Relay                   MAJ
         Notifies                Enabled
         Syslog                  Enabled
Temperature-Secondary
         Alarm                   Disabled
         Threshold
         Relay
         Notifies                Disabled
         Syslog                  Disabled
SD-Card
         Alarm                   Disabled
         Relay
         Notifies                Disabled
         Syslog                  Enabled
Input-Alarm 1
```

```
          Alarm                   Enabled
          Relay
          Notifies                Disabled
          Syslog                  Enabled
Input-Alarm 2
          Alarm                   Enabled
          Relay
          Notifies                Disabled
          Syslog                  Enabled
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
Alarm input and output ports.	*Cisco IE 2000 Switch Hardware Installation Guide*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 4

Performing Switch Setup Configuration

Restrictions for Performing Switch Setup Configuration

- The DHCP-based autoconfiguration with a saved configuration process stops if there is not at least one Layer 3 interface in an up state without an assigned IP address in the network.
- Unless you configure a timeout, the DHCP-based autoconfiguration with a saved configuration feature tries indefinitely to download an IP address.
- The auto-install process stops if a configuration file cannot be downloaded or it the configuration file is corrupted.

Note The configuration file that is downloaded from TFTP is merged with the existing configuration in the running configuration but is not saved in the NVRAM unless you enter the **write memory** or **copy running-configuration startup-configuration** privileged EXEC command. Note that if the downloaded configuration is saved to the startup configuration, the feature is not triggered during subsequent system restarts.

Information About Performing Switch Setup Configuration

This chapter describes how to perform your initial switch configuration tasks that include IP address assignments and DHCP autoconfiguration.

Switch Boot Process

To start your switch, you need to follow the procedures in the *Cisco IE 2000 Switch Getting Started Guide* or the hardware installation guide for installing and powering on the switch and for setting up the initial switch configuration (IP address, subnet mask, default gateway, secret and Telnet passwords, and so forth).

The normal boot process involves the operation of the boot loader software, which performs these activities:

- Performs low-level CPU initialization—Initializes the CPU registers, which control where physical memory is mapped, its quantity and its speed.
- Performs power-on self-test (POST) for the CPU subsystem—Tests the CPU DRAM and the portion of the flash device that makes up the flash file system.

- Initializes the flash memory card file system on the system board.
- Loads a default operating system software image into memory and boots up the switch.

The boot loader provides access to the flash file system before the operating system is loaded. Normally, the boot loader is used only to load, uncompress, and launch the operating system. After the boot loader gives the operating system control of the CPU, the boot loader is not active until the next system reset or power-on.

The switch supports a flash memory card that makes it possible to replace a failed switch without reconfiguring the new switch. The slot for the flash memory card is hot swappable and front-accessed. A cover protects the flash card and holds the card firmly in place. The cover is hinged and closed with a captive screw, which prevents the card from coming loose and protects against shock and vibration.

Use the **show flash:** privileged EXEC command to display the flash memory card file settings. For information about how to remove or replace the flash memory card on the switch, see the *Cisco IE 2000 Hardware Installation Guide*.

The boot loader also provides trap-door access into the system if the operating system has problems serious enough that it cannot be used. The trap-door mechanism provides enough access to the system so that if it is necessary, you can format the flash file system, reinstall the operating system software image by using the Xmodem Protocol, recover from a lost or forgotten password, and finally restart the operating system. For more information, see "Recovering from Software Failures" and the "Recovering from a Lost or Forgotten Password".

Note You can disable password recovery. For more information, see "Disabling Password Recovery".

Before you can assign switch information, make sure you have connected a PC or terminal to the console port, and configured the PC or terminal-emulation software baud rate and character format to match these of the switch console port:

- Baud rate default is 9600.
- Data bits default is 8.

  Note If the data bits option is set to 8, set the parity option to none.

- Stop bits default is 1.
- Parity settings default is none.

Default Switch Boot Settings

Feature	Default Setting
Operating system software image	The switch attempts to automatically boot up the system using information in the BOOT environment variable. If the variable is not set, the switch attempts to load and execute the first executable image it can by performing a recursive, depth-first search throughout the flash file system. The Cisco IOS image is stored in a directory that has the same name as the image file (excluding the .bin extension). In a depth-first search of a directory, each encountered subdirectory is completely searched before continuing the search in the original directory.
Configuration file	Configured switches use the *config.text* file stored on the system board in flash memory. A new switch has no configuration file.

Switch Boot Optimization

The normal switch boot process involves a memory test, file system check (FSCK), and power-on self-test (POST).

The **boot fast** command in global configuration mode is enabled by default to permit switch boot optimization, which disables these tests and minimizes the bootup time. However, after a system crash this feature is automatically disabled.

Reload sequences occur immediately if your switch is set up to automatically bring up the system by using information in the BOOT environment variable. Otherwise, these reload sequences occur after you enter the manual **boot** command in bootloader configuration mode.

First Reload

The switch disables the boot fast feature and displays the following warning message:

```
"Reloading with boot fast feature disabled"
```

After the system message appears, the system saves the crash information and automatically resets itself for the next reload cycle.

Second Reload

The boot loader performs its normal full memory test and FSCK check with LED status progress. If the memory and FSCK tests are successful, the system performs additional POST tests and the results are displayed on the console.

The boot fast feature is reenabled after the system comes up successfully.

Switch Information Assignment

You can assign IP information through the switch setup program, through a DHCP server, or manually.

Use the switch setup program if you want to be prompted for specific IP information. With this program, you can also configure a hostname and an enable secret password. The program gives you the option of assigning a Telnet password (to provide security during remote management) and configuring your switch as a command or member switch of a cluster or as a standalone switch. For more information about the setup program, see the hardware installation guide.

Use a DHCP server for centralized control and automatic assignment of IP information after the server is configured.

Note If you are using DHCP, do not respond to any of the questions in the setup program until the switch receives the dynamically assigned IP address and reads the configuration file.

If you are an experienced user familiar with the switch configuration steps, manually configure the switch. Otherwise, use the setup program.

Switch Default Settings

Table 4-1 Switch Default Settings

Feature	Default Setting
IP address and subnet mask	No IP address or subnet mask is defined.
Default gateway	No default gateway is defined.
Enable secret password	No password is defined.
Hostname	The factory-assigned default hostname is *Switch*.
Telnet password	No password is defined.
Cluster command switch functionality	Disabled.
Cluster name	No cluster name is defined.
Manual boot	No.
Boot optimization	Enabled.

DHCP-Based Autoconfiguration Overview

DHCP provides configuration information to Internet hosts and internetworking devices. This protocol consists of two components: one for delivering configuration parameters from a DHCP server to a device and a mechanism for allocating network addresses to devices. DHCP is built on a client-server model, in which designated DHCP servers allocate network addresses and deliver configuration parameters to dynamically configured devices. The switch can act as both a DHCP client and a DHCP server.

During DHCP-based autoconfiguration, your switch (DHCP client) is automatically configured at startup with IP address information and a configuration file.

With DHCP-based autoconfiguration, no DHCP client-side configuration is needed on your switch. However, you need to configure the DHCP server for various lease options associated with IP addresses. If you are using DHCP to relay the configuration file location on the network, you might also need to configure a Trivial File Transfer Protocol (TFTP) server and a Domain Name System (DNS) server.

The DHCP server for your switch can be on the same LAN or on a different LAN than the switch. If the DHCP server is running on a different LAN, you should configure a DHCP relay device between your switch and the DHCP server. A relay device forwards broadcast traffic between two directly connected LANs. A router does not forward broadcast packets, but it forwards packets based on the destination IP address in the received packet.

DHCP-based autoconfiguration replaces the BOOTP client functionality on your switch.

DHCP Client Request Process

When you boot up your switch, the DHCP client is invoked and requests configuration information from a DHCP server when the configuration file is not present on the switch. If the configuration file is present and the configuration includes the **ip address dhcp** interface configuration command on specific routed interfaces, the DHCP client is invoked and requests the IP address information for those interfaces.

Figure 4-1 shows the sequence of messages that are exchanged between the DHCP client and the DHCP server.

Figure 4-1 DHCP Client and Server Message Exchange

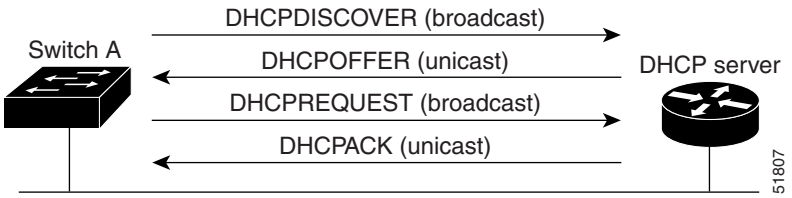

The client, Switch A, broadcasts a DHCPDISCOVER message to locate a DHCP server. The DHCP server offers configuration parameters (such as an IP address, subnet mask, gateway IP address, DNS IP address, a lease for the IP address, and so forth) to the client in a DHCPOFFER unicast message.

In a DHCPREQUEST broadcast message, the client returns a formal request for the offered configuration information to the DHCP server. The formal request is broadcast so that all other DHCP servers that received the DHCPDISCOVER broadcast message from the client can reclaim the IP addresses that they offered to the client.

The DHCP server confirms that the IP address has been allocated to the client by returning a DHCPACK unicast message to the client. With this message, the client and server are bound, and the client uses configuration information received from the server. The amount of information the switch receives depends on how you configure the DHCP serverd in conjunction with the TFTP server. For more information, see the "TFTP Server" section on page 4-7.

If the configuration parameters sent to the client in the DHCPOFFER unicast message are invalid (a configuration error exists), the client returns a DHCPDECLINE broadcast message to the DHCP server.

The DHCP server sends the client a DHCPNAK denial broadcast message, which means that the offered configuration parameters have not been assigned, that an error has occurred during the negotiation of the parameters, or that the client has been slow in responding to the DHCPOFFER message. (The DHCP server assigned the parameters to another client.)

A DHCP client might receive offers from multiple DHCP or BOOTP servers and can accept any of the offers; however, the client usually accepts the first offer it receives. The offer from the DHCP server is not a guarantee that the IP address is allocated to the switch. However, the server usually reserves the address until the client has had a chance to formally request the address. If the switch accepts replies from a BOOTP server and configures itself, the switch broadcasts, instead of unicasts, TFTP requests to obtain the switch configuration file.

The DHCP hostname option allows a group of switches to obtain hostnames and a standard configuration from the central management DHCP server. A client (switch) includes in its DCHPDISCOVER message an option 12 field used to request a hostname and other configuration parameters from the DHCP server. The configuration files on all clients are identical except for their DHCP-obtained hostnames.

If a client has a default hostname (the **hostname** *name* global configuration command is not configured or the **no hostname** global configuration command is entered to remove the hostname), the DHCP hostname option is not included in the packet when you enter the **ip address dhcp** interface configuration command. In this case, if the client receives the DCHP hostname option from the DHCP interaction while acquiring an IP address for an interface, the client accepts the DHCP hostname option and sets the flag to show that the system now has a hostname configured.

DHCP-Based Autoconfiguration and Image Update

You can use the DHCP image upgrade features to configure a DHCP server to download both a new image and a new configuration file to one or more switches in a network. This helps ensure that each new switch added to a network receives the same image and configuration.

There are two types of DHCP image upgrades: DHCP autoconfiguration and DHCP auto-image update.

DHCP Autoconfiguration

DHCP autoconfiguration downloads a configuration file to one or more switches in your network from a DHCP server. The downloaded configuration file becomes the running configuration of the switch. It does not over write the bootup configuration saved in the flash, until you reload the switch.

DHCP Auto-Image Update

You can use DHCP auto-image upgrade with DHCP autoconfiguration to download both a configuration *and* a new image to one or more switches in your network. The switch (or switches) downloading the new configuration and the new image can be blank (or only have a default factory configuration loaded).

If the new configuration is downloaded to a switch that already has a configuration, the downloaded configuration is appended to the configuration file stored on the switch. (Any existing configuration is not overwritten by the downloaded one.)

Note To enable a DHCP auto-image update on the switch, the TFTP server where the image and configuration files are located must be configured with the correct option 67 (the configuration filename), option 66 (the DHCP server hostname) option 150 (the TFTP server address), and option 125 (description of the file) settings.

For procedures to configure the switch as a DHCP server, see the "DHCP Server Configuration Guidelines" section on page 4-7 and the "Configuring DHCP" section of the "IP addressing and Services" section of the *Cisco IOS IP DHCP Configuration Guide*, Release 15.0.

After you install the switch in your network, the auto-image update feature starts. The downloaded configuration file is saved in the running configuration of the switch, and the new image is downloaded and installed on the switch. When you reboot the switch, the configuration is stored in the saved configuration on the switch.

DHCP Server Configuration Guidelines

Follow these guidelines if you are configuring a device as a DHCP server:

- Configure the DHCP server with reserved leases that are bound to each switch by the switch hardware address.
- If you want the switch to receive IP address information, you must configure the DHCP server with these lease options:
  - IP address of the client (required)
  - Subnet mask of the client (required)
  - Router IP address (default gateway address to be used by the switch) (required)
  - DNS server IP address (optional)
- If you want the switch to receive the configuration file from a TFTP server, you must configure the DHCP server with these lease options:
  - TFTP server name (required)
  - Boot filename (the name of the configuration file that the client needs) (recommended)
  - Hostname (optional)
- Depending on the settings of the DHCP server, the switch can receive IP address information, the configuration file, or both.
- If you do not configure the DHCP server with the lease options described previously, it replies to client requests with only those parameters that are configured.

  If the IP address and the subnet mask are not in the reply, the switch is not configured. If the router IP address or the TFTP server name are not found, the switch might send broadcast, instead of unicast, TFTP requests. Unavailability of other lease options does not affect autoconfiguration.
- The switch can act as a DHCP server. By default, the Cisco IOS DHCP server and relay agent features are enabled on your switch but are not configured. These features are not operational. If your DHCP server is a Cisco device, for additional information about configuring DHCP, see the "Configuring DHCP" section of the "IP Addressing and Services" section of the *Cisco IOS IP Configuration Guide* on Cisco.com.

TFTP Server

Based on the DHCP server configuration, the switch attempts to download one or more configuration files from the TFTP server. If you configured the DHCP server to respond to the switch with all the options required for IP connectivity to the TFTP server, and if you configured the DHCP server with a TFTP server name, address, and configuration filename, the switch attempts to download the specified configuration file from the specified TFTP server.

If you did not specify the configuration filename, the TFTP server, or if the configuration file could not be downloaded, the switch attempts to download a configuration file by using various combinations of filenames and TFTP server addresses. The files include the specified configuration filename (if any) and

these files: network-config, cisconet.cfg, and *hostname*.config (or *hostname*.cfg), where *hostname* is the switch's current hostname. The TFTP server addresses used include the specified TFTP server address (if any) and the broadcast address (255.255.255.255).

For the switch to successfully download a configuration file, the TFTP server must contain one or more configuration files in its base directory. The files can include these files:

- The configuration file named in the DHCP reply (the actual switch configuration file).
- The network-confg or the cisconet.cfg file (known as the default configuration files).
- The router-confg or the ciscortr.cfg file (These files contain commands common to all switches. Normally, if the DHCP and TFTP servers are properly configured, these files are not accessed.)

If you specify the TFTP server name in the DHCP server-lease database, you must also configure the TFTP server name-to-IP-address mapping in the DNS-server database.

If the TFTP server to be used is on a different LAN from the switch, or if it is to be accessed by the switch through the broadcast address (which occurs if the DHCP server response does not contain all the required information described previously), a relay must be configured to forward the TFTP packets to the TFTP server. For more information, see the "Relay Device" section on page 4-8. The preferred solution is to configure the DHCP server with all the required information.

DNS Server

The DHCP server uses the DNS server to resolve the TFTP server name to an IP address. You must configure the TFTP server name-to-IP address map on the DNS server. The TFTP server contains the configuration files for the switch.

You can configure the IP addresses of the DNS servers in the lease database of the DHCP server from where the DHCP replies will retrieve them. You can enter up to two DNS server IP addresses in the lease database.

The DNS server can be on the same or on a different LAN as the switch. If it is on a different LAN, the switch must be able to access it through a router.

Relay Device

You must configure a relay device, also referred to as a *relay agent*, when a switch sends broadcast packets that require a response from a host on a different LAN. Examples of broadcast packets that the switch might send are DHCP, DNS, and in some cases, TFTP packets. You must configure this relay device to forward received broadcast packets on an interface to the destination host.

If the relay device is a Cisco router, enable IP routing (**ip routing** global configuration command), and configure helper addresses by using the **ip helper-address** interface configuration command.

For example, in Figure 4-2, configure the router interfaces as follows:

On interface 10.0.0.2:

```
router(config-if)# ip helper-address 20.0.0.2
router(config-if)# ip helper-address 20.0.0.3
router(config-if)# ip helper-address 20.0.0.4
```

On interface 20.0.0.1:

```
router(config-if)# ip helper-address 10.0.0.1
```

Figure 4-2 Relay Device Used in Autoconfiguration

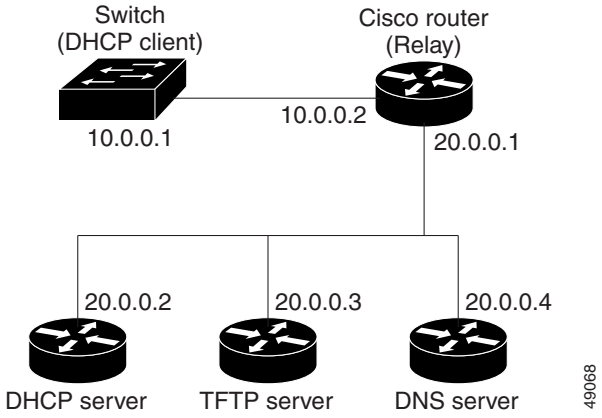

How to Obtain Configuration Files

Depending on the availability of the IP address and the configuration filename in the DHCP reserved lease, the switch obtains its configuration information in these ways:

- The IP address and the configuration filename is reserved for the switch and provided in the DHCP reply (one-file read method).

  The switch receives its IP address, subnet mask, TFTP server address, and the configuration filename from the DHCP server. The switch sends a unicast message to the TFTP server to retrieve the named configuration file from the base directory of the server and upon receipt, it completes its boot-up process.

- The IP address and the configuration filename is reserved for the switch, but the TFTP server address is not provided in the DHCP reply (one-file read method).

  The switch receives its IP address, subnet mask, and the configuration filename from the DHCP server. The switch sends a broadcast message to a TFTP server to retrieve the named configuration file from the base directory of the server, and upon receipt, it completes its boot-up process.

- Only the IP address is reserved for the switch and provided in the DHCP reply. The configuration filename is not provided (two-file read method).

  The switch receives its IP address, subnet mask, and the TFTP server address from the DHCP server. The switch sends a unicast message to the TFTP server to retrieve the network-confg or cisconet.cfg default configuration file. (If the network-confg file cannot be read, the switch reads the cisconet.cfg file.)

  The default configuration file contains the hostnames-to-IP-address mapping for the switch. The switch fills its host table with the information in the file and obtains its hostname. If the hostname is not found in the file, the switch uses the hostname in the DHCP reply. If the hostname is not specified in the DHCP reply, the switch uses the default *Switch* as its hostname.

  After obtaining its hostname from the default configuration file or the DHCP reply, the switch reads the configuration file that has the same name as its hostname (*hostname*-confg or *hostname*.cfg, depending on whether network-confg or cisconet.cfg was read earlier) from the TFTP server. If the cisconet.cfg file is read, the filename of the host is truncated to eight characters.

  If the switch cannot read the network-confg, cisconet.cfg, or the hostname file, it reads the router-confg file. If the switch cannot read the router-confg file, it reads the ciscortr.cfg file.

Note The switch broadcasts TFTP server requests if the TFTP server is not obtained from the DHCP replies, if all attempts to read the configuration file through unicast transmissions fail, or if the TFTP server name cannot be resolved to an IP address.

How to Control Environment Variables

With a normally operating switch, you enter the boot loader mode only through a switch console connection configured for 9600 b/s. Unplug the switch power cord, and press the switch **Mode** button while reconnecting the power cord. You can release the **Mode** button a second or two after the LED above port 1 turns off. Then the boot loader *switch:* prompt appears.

The switch boot loader software provides support for nonvolatile environment variables, which can be used to control how the boot loader or any other software running on the system behaves. Boot loader environment variables are similar to environment variables that can be set on UNIX or DOS systems.

Environment variables that have values are stored in flash memory outside of the flash file system.

Each line in these files contains an environment variable name and an equal sign followed by the value of the variable. A variable has no value if it is not listed in this file; it has a value if it is listed in the file even if the value is a null string. A variable that is set to a null string (for example, " ") is a variable with a value. Many environment variables are predefined and have default values.

Environment variables store two kinds of data:

- Data that controls code, which does not read the Cisco IOS configuration file. For example, the name of a boot loader helper file, which extends or patches the functionality of the boot loader can be stored as an environment variable.
- Data that controls code, which is responsible for reading the Cisco IOS configuration file. For example, the name of the Cisco IOS configuration file can be stored as an environment variable.

You can change the settings of the environment variables by accessing the boot loader or by using Cisco IOS commands. Under normal circumstances, it is not necessary to alter the setting of the environment variables.

Note For complete syntax and usage information for the boot loader commands and environment variables, see the command reference for this release.

Common Environment Variables

Table 4-2 describes the function of the most common environment variables.

Table 4-2 Environment Variables

Variable	Boot Loader Command	Cisco IOS Global Configuration Command
BOOT	**set BOOT** *filesystem***:/***file-url* ... A semicolon-separated list of executable files to try to load and execute when automatically booting. If the BOOT environment variable is not set, the system attempts to load and execute the first executable image it can find by using a recursive, depth-first search through the flash file system. If the BOOT variable is set but the specified images cannot be loaded, the system attempts to boot the first bootable file that it can find in the flash file system.	**boot system** f*ilesystem:/file-url* ... Specifies the Cisco IOS image to load during the next boot cycle. This command changes the setting of the BOOT environment variable.
MANUAL_BOOT	**set MANUAL_BOOT yes** Decides whether the switch automatically or manually boots up. Valid values are 1, yes, 0, and no. If it is set to no or 0, the boot loader attempts to automatically boot up the system. If it is set to anything else, you must manually boot up the switch from the boot loader mode.	**boot manual** Enables manually booting up the switch during the next boot cycle and changes the setting of the MANUAL_BOOT environment variable. The next time you reboot the system, the switch is in boot loader mode. To boot up the system, use the **boot flash:***filesystem***:/***file-url* boot loader command, and specify the name of the bootable image.
CONFIG_FILE	**set CONFIG_FILE flash:/***file-url* Changes the filename that Cisco IOS uses to read and write a nonvolatile copy of the system configuration.	**boot config-file flash:/***file-url* Specifies the filename that Cisco IOS uses to read and write a nonvolatile copy of the system configuration. This command changes the CONFIG_FILE environment variable.

Scheduled Reload of the Software Image

You can schedule a reload of the software image to occur on the switch at a later time (for example, late at night or during the weekend when the switch is used less), or you can synchronize a reload network-wide (for example, to perform a software upgrade on all switches in the network).

Note A scheduled reload must take place within approximately 24 days.

You have these reload options:

- Software reload to take effect in the specified minutes or hours and minutes. The reload must take place within approximately 24 days. You can specify the reason for the reload in a string up to 255 characters in length.
- Software reload to take place at the specified time (using a 24-hour clock). If you specify the month and day, the reload is scheduled to take place at the specified time and date. If you do not specify the month and day, the reload takes place at the specified time on the current day (if the specified time is later than the current time) or on the next day (if the specified time is earlier than the current time). Specifying 00:00 schedules the reload for midnight.

The **reload** command halts the system. If the system is not set to manually boot up, it reboots itself.

If your switch is configured for manual booting, do not reload it from a virtual terminal. This restriction prevents the switch from entering the boot loader mode and thereby taking it from the remote user's control.

If you modify your configuration file, the switch prompts you to save the configuration before reloading. During the save operation, the system requests whether you want to proceed with the save if the CONFIG_FILE environment variable points to a startup configuration file that no longer exists. If you proceed in this situation, the system enters setup mode upon reload.

To cancel a previously scheduled reload, use the **reload cancel** privileged EXEC command.

How to Perform Switch Setup Configuration

Using DHCP to download a new image and a new configuration to a switch requires that you configure at least two switches. One switch acts as a DHCP and TFTP server and the second switch (client) is configured to download either a new configuration file or a new configuration file and a new image file.

Configuring DHCP Autoconfiguration (Only Configuration File)

This task describes how to configure DHCP autoconfiguration of the TFTP and DHCP settings on a new switch to download a new configuration file.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip dhcp** *poolname*	Creates a name for the DHCP Server address pool, and enters DHCP pool configuration mode.
Step 3	**bootfile** *filename*	Specifies the name of the configuration file that is used as a boot image.
Step 4	**network** *network-number mask prefix-length*	Specifies the subnet network number and mask of the DHCP address pool. **Note** The prefix length specifies the number of bits that comprise the address prefix. The prefix is an alternative way of specifying the network mask of the client. The prefix length must be preceded by a forward slash (/).
Step 5	**default-router** *address*	Specifies the IP address of the default router for a DHCP client.
Step 6	**option 150** *address*	Specifies the IP address of the TFTP server.

	Command	Purpose
Step 7	**exit**	Returns to global configuration mode.
Step 8	**tftp-server flash**:*filename.text*	Specifies the configuration file on the TFTP server.
Step 9	**interface** *interface-id*	Specifies the address of the client that will receive the configuration file.
Step 10	**no switchport**	Puts the interface into Layer 3 mode.
Step 11	**ip address** *address mask*	Specifies the IP address and mask for the interface.
Step 12	**end**	Returns to privileged EXEC mode.
Step 13	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring DHCP Auto-Image Update (Configuration File and Image)

This task describes DHCP autoconfiguration to configure TFTP and DHCP settings on a new switch to download a new image and a new configuration file.

Before You Begin

You must create a text file (for example, autoinstall_dhcp) that will be uploaded to the switch. In the text file, put the name of the image that you want to download. This image must be a tar and not a bin file.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip dhcp pool** *name*	Creates a name for the DHCP server address pool and enters DHCP pool configuration mode.
Step 3	**bootfile** *filename*	Specifies the name of the file that is used as a boot image.
Step 4	**network** *network-number mask prefix-length*	Specifies the subnet network number and mask of the DHCP address pool. **Note** The prefix length specifies the number of bits that comprise the address prefix. The prefix is an alternative way of specifying the network mask of the client. The prefix length must be preceded by a forward slash (/).
Step 5	**default-router** *address*	Specifies the IP address of the default router for a DHCP client.
Step 6	**option 150** *address*	Specifies the IP address of the TFTP server.
Step 7	**option 125** *hex*	Specifies the path to the text file that describes the path to the image file.
Step 8	**copy tftp flash** *filename.txt*	Uploads the text file to the switch.
Step 9	**copy tftp flash** *imagename.tar*	Uploads the tar file for the new image to the switch.
Step 10	**exit**	Returns to global configuration mode.
Step 11	**tftp-server flash:***config.text*	Specifies the Cisco IOS configuration file on the TFTP server.
Step 12	**tftp-server flash:***imagename.tar*	Specifies the image name on the TFTP server.
Step 13	**tftp-server flash:***filename.txt*	Specifies the text file that contains the name of the image file to download.
Step 14	**interface** *interface-id*	Specifies the address of the client that will receive the configuration file.
Step 15	**no switchport**	Puts the interface into Layer 3 mode.

	Command	Purpose
Step 16	**ip address** *address mask*	Specifies the IP address and mask for the interface.
Step 17	**end**	Returns to privileged EXEC mode.
Step 18	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring the Client

You should only configure and enable the Layer 3 interface. Do not assign an IP address or DHCP-based autoconfiguration with a saved configuration.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**boot host dhcp**	Enables autoconfiguration with a saved configuration.
Step 3	**boot host retry timeout** *timeout-value*	(Optional) Sets the amount of time the system tries to download a configuration file. **Note** If you do not set a timeout, the system tries indefinitely to obtain an IP address from the DHCP server.
Step 4	**banner config-save ^C** *warning-message* **^C**	(Optional) Creates warning messages to be displayed when you try to save the configuration file to NVRAM.
Step 5	**end**	Returns to privileged EXEC mode.
Step 6	**show boot**	Verifies the configuration.

Manually Assigning IP Information on a Routed Port

This task describes how to manually assign IP information on a Layer 3 routed port.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *type id*	Enters interface configuration mode.
Step 3	**no switchport**	Configures an interface into Layer 3 mode.
Step 4	**ip address** *address mask*	Specifies the IP address and mask for the interface.
Step 5	**exit**	Returns to global configuration mode.
Step 6	**ip default-gateway** *ip-address*	Enters the IP address of the next-hop router interface that is directly connected to the switch where a default gateway is being configured. The default gateway receives IP packets with unresolved destination IP addresses from the switch. Once the default gateway is configured, the switch has connectivity to the remote networks with which a host needs to communicate. **Note** When your switch is configured to route with IP, it does not need to have a default gateway set.
Step 7	**end**	Returns to privileged EXEC mode.

	Command	Purpose
Step 8	**show ip redirects**	Verifies the configured default gateway.
Step 9	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Manually Assigning IP Information to SVIs

This task describes how to manually assign IP information to multiple switched virtual interfaces (SVIs).

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface vlan** *vlan-id*	Enters interface configuration mode, and enters the VLAN to which the IP information is assigned. The VLAN range is 1 to 4096.
Step 3	**ip address** *ip-address subnet-mask*	Enters the IP address and subnet mask.
Step 4	**exit**	Returns to global configuration mode.
Step 5	**ip default-gateway** *ip-address*	Enters the IP address of the next-hop router interface that is directly connected to the switch where a default gateway is being configured. The default gateway receives IP packets with unresolved destination IP addresses from the switch. Once the default gateway is configured, the switch has connectivity to the remote networks with which a host needs to communicate. **Note** When your switch is configured to route with IP, it does not need to have a default gateway set.
Step 6	**end**	Returns to privileged EXEC mode.
Step 7	**show interfaces vlan** *vlan-id*	Verifies the configured IP address.
Step 8	**show ip redirects**	Verifies the configured default gateway.
Step 9	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Modifying the Startup Configuration

Specifying the Filename to Read and Write the System Configuration

By default, the Cisco IOS software uses the *config.text* file to read and write a nonvolatile copy of the system configuration. However, you can specify a different filename, which will be loaded during the next boot-up cycle.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**boot config-file flash:***/file-url*	Specifies the configuration file to load during the next boot-up cycle. For *file-url*, specify the path (directory) and the configuration filename. Filenames and directory names are case sensitive.
Step 3	**end**	Returns to privileged EXEC mode.
Step 4	**show boot**	Verifies your entries. The **boot config-file** global configuration command changes the setting of the CONFIG_FILE environment variable.
Step 5	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Manually Booting the Switch

By default, the switch automatically boots up; however, you can configure it to manually boot up.

Before You Begin

Use a standalone switch for this task.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**boot manual**	Enables the switch to manually boot up during the next boot cycle.
Step 3	**end**	Returns to privileged EXEC mode.
Step 4	**show boot**	Verifies your entries. The **boot manual** global command changes the setting of the MANUAL_BOOT environment variable. The next time you reboot the system, the switch is in boot loader mode, shown by the switch: prompt. To boot up the system, use the **boot** *filesystem***:***/file-url* boot loader command. • For *filesystem***:**, use **flash:** for the system board flash device. • For *file-url*, specify the path (directory) and the name of the bootable image. Filenames and directory names are case sensitive.
Step 5	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Booting a Specific Software Image

By default, the switch attempts to automatically boot up the system using information in the BOOT environment variable. If this variable is not set, the switch attempts to load and execute the first executable image it can by performing a recursive, depth-first search throughout the flash file system. In a depth-first search of a directory, each encountered subdirectory is completely searched before continuing the search in the original directory. However, you can specify a specific image to boot up.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**boot system** *filesystem***:/***file-url*	Configures the switch to boot a specific image in flash memory during the next boot cycle. • For *filesystem***:**, use **flash:** for the system board flash device. • For *file-url*, specify the path (directory) and the name of the bootable image. Filenames and directory names are case sensitive.
Step 3	**end**	Returns to privileged EXEC mode.
Step 4	**show boot**	Verifies your entries. The **boot system** global command changes the setting of the BOOT environment variable. During the next boot cycle, the switch attempts to automatically boot up the system using information in the BOOT environment variable.
Step 5	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Monitoring Switch Setup Configuration

Verifying the Switch Running Configuration

You can check the configuration settings that you entered or changes that you made by entering this privileged EXEC command:

```
Switch# show running-config
Building configuration...

Current configuration: 1363 bytes
!
version 12.2
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname Switch A
!
enable secret 5 $1$ej9.$DMUvAUnZOAmvmgqBEzIxE0
!
.
<output truncated>
```

```
.
interface gigabitethernet1/1
no switchport
ip address 172.20.137.50 255.255.255.0
!
interface gigabitethernet1/2
mvr type source

<output truncated>

...!
interface VLAN1
 ip address 172.20.137.50 255.255.255.0
 no ip directed-broadcast
!
ip default-gateway 172.20.137.1 !
!
snmp-server community private RW
snmp-server community public RO
snmp-server community private@es0 RW
snmp-server community public@es0 RO
snmp-server chassis-id 0x12
!
end
```

To store the configuration or changes you have made to your startup configuration in flash memory, enter this privileged EXEC command:

```
Switch# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
```

This command saves the configuration settings that you made. If you fail to do this, your configuration will be lost the next time you reload the system. To display information stored in the NVRAM section of flash memory, use the **show startup-config** or **more startup-config** privileged EXEC command.

For more information about alternative locations from which to copy the configuration file, see Appendix A, "Working with the Cisco IOS File System, Configuration Files, and Software Images."

Configuration Examples for Performing Switch Setup Configuration

Retrieving IP Information Using DHCP-Based Autoconfiguration: Example

Switch A reads its configuration file as follows:

- It obtains its IP address 10.0.0.21 from the DHCP server.
- If no configuration filename is given in the DHCP server reply, Switch A reads the network-confg file from the base directory of the TFTP server.
- It adds the contents of the network-confg file to its host table.
- It reads its host table by indexing its IP address 10.0.0.21 to its hostname (switcha).
- It reads the configuration file that corresponds to its hostname; for example, it reads *switch1-confg* from the TFTP server.

Switches B through D retrieve their configuration files and IP addresses in the same way.

Figure 4-3 shows a sample network for retrieving IP information by using DHCP-based autoconfiguration.

Figure 4-3 DHCP-Based Autoconfiguration Network Example

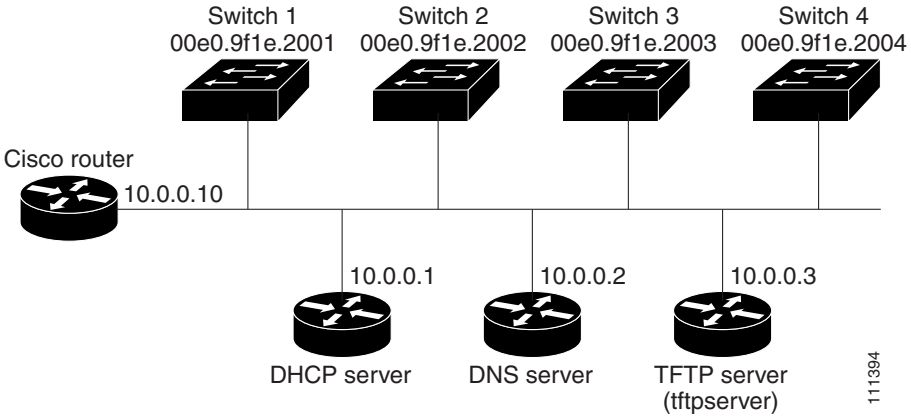

Table 4-3 shows the configuration of the reserved leases on the DHCP server.

Table 4-3 DHCP Server Configuration

	Switch A	Switch B	Switch C	Switch D
Binding key (hardware address)	00e0.9f1e.2001	00e0.9f1e.2002	00e0.9f1e.2003	00e0.9f1e.2004
IP address	10.0.0.21	10.0.0.22	10.0.0.23	10.0.0.24
Subnet mask	255.255.255.0	255.255.255.0	255.255.255.0	255.255.255.0
Router address	10.0.0.10	10.0.0.10	10.0.0.10	10.0.0.10
DNS server address	10.0.0.2	10.0.0.2	10.0.0.2	10.0.0.2
TFTP server name	*tftpserver* or *10.0.0.3*	*tftpserver* or *10.0.0.3*	*tftpserver* or *10.0.0.3*	*tftpserver* or *10.0.0.3*
Boot filename (configuration file) (optional)	switcha-confg	switchb-confg	switchc-confg	switchd-confg
Hostname (optional)	switcha	switchb	switchc	switchd

DNS Server Configuration

The DNS server maps the TFTP server name *tftpserver* to IP address 10.0.0.3.

TFTP Server Configuration (on UNIX)

The TFTP server base directory is set to /tftpserver/work/. This directory contains the network-confg file used in the two-file read method. This file contains the hostname to be assigned to the switch based on its IP address. The base directory also contains a configuration file for each switch (*switcha-confg*, *switchb-confg*, and so forth) as shown in this display:

```
prompt> cd /tftpserver/work/
prompt> ls
network-confg
switcha-confg
switchb-confg
switchc-confg
```

```
switchd-confg
prompt> cat network-confg
ip host switcha 10.0.0.21
ip host switchb 10.0.0.22
ip host switchc 10.0.0.23
ip host switchd 10.0.0.24
```

DHCP Client Configuration

No configuration file is present on Switch A through Switch D.

Scheduling Software Image Reload: Examples

This example shows how to reload the software on the switch on the current day at 7:30 p.m:

```
Switch# reload at 19:30
Reload scheduled for 19:30:00 UTC Wed Jun 5 1996 (in 2 hours and 25 minutes)
Proceed with reload? [confirm]
```

This example shows how to reload the software on the switch at a future time:

```
Switch# reload at 02:00 jun 20
Reload scheduled for 02:00:00 UTC Thu Jun 20 1996 (in 344 hours and 53 minutes)
Proceed with reload? [confirm]
```

To cancel a previously scheduled reload, use the **reload cancel** privileged EXEC command.

Configuring DHCP Auto-Image Update: Example

```
Switch# configure terminal
Switch(config)# ip dhcp pool pool1
Switch(dhcp-config)# network 10.10.10.0 255.255.255.0
Switch(dhcp-config)# bootfile config-boot.text
Switch(dhcp-config)# default-router 10.10.10.1
Switch(dhcp-config)# option 150 10.10.10.1
Switch(dhcp-config)# exit
Switch(config)# tftp-server flash:config-boot.text
Switch(config)# interface gigabitethernet1/2
Switch(config-if)# no switchport
Switch(config-if)# ip address 10.10.10.1 255.255.255.0
Switch(config-if)# end
```

Configuring a Switch as a DHCP Server: Example

This example shows how to configure a switch as a DHCP server so it downloads a configuration file:

```
Switch# config terminal
Switch(config)# ip dhcp pool pool1
Switch(dhcp-config)# network 10.10.10.0 255.255.255.0
Switch(dhcp-config)# bootfile config-boot.text
Switch(dhcp-config)# default-router 10.10.10.1
Switch(dhcp-config)# option 150 10.10.10.1
Switch(dhcp-config)# option 125 hex
0000.0009.0a05.08661.7574.6f69.6e73.7461.6c6c.5f64.686370
Switch(dhcp-config)# exit
Switch(config)# tftp-server flash:config-boot.text
Switch(config)# tftp-server flash:c-ipservices-mz.122-44.3.SE.tar
Switch(config)# tftp-server flash:ies-lanbase-tar.122-44.EX.tar
```

```
Switch(config)# tftp-server flash:boot-config.text
Switch(config)# tftp-server flash: autoinstall_dhcp
Switch(config)# interface gigabitethernet1/2
Switch(config-if)# no switchport
Switch(config-if)# ip address 10.10.10.1 255.255.255.0
Switch(config-if)# end
```

Configuring Client to Download Files from DHCP Server

This example uses a Layer 3 SVI interface on VLAN 99 to enable DHCP-based autoconfiguration with a saved configuration:

```
Switch# configure terminal
Switch(conf)# boot host dhcp
Switch(conf)# boot host retry timeout 300
Switch(conf)# banner config-save ^C Caution - Saving Configuration File to NVRAM May Cause
You to Nolonger Automatically Download Configuration Files at Reboot^C
Switch(config)# vlan 99
Switch(config-vlan)# interface vlan 99
Switch(config-if)# no shutdown
Switch(config-if)# end
Switch# show boot
BOOT path-list:
Config file:          flash:/config.text
Private Config file:  flash:/private-config.text
Enable Break:         no
Manual Boot:          no
HELPER path-list:
NVRAM/Config file
      buffer size:    32768
Timeout for Config
          Download:    300 seconds
Config Download
       via DHCP:       enabled (next boot: enabled)
Switch#
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 5

Configuring Cisco IOS Configuration Engine

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for Configuring Cisco IOS Configuration Engine

Set the CNS DeviceID

- When using the Cisco Configuration Engine user interface, you must first set the DeviceID field to the hostname value that the switch acquires after, not before, you use the **cns config initial** global configuration command at the switch. Otherwise, subsequent **cns config partial** global configuration command operations malfunction.

Enable Automated CNS Configuration

- To enable automated CNS configuration of the switch, you must first complete the prerequisites in Table 5-1. When you complete them, power on the switch. At the **setup** prompt, you do not need to enter a command. The switch begins the initial configuration as described in the "Initial Configuration" section on page 5-5. When the full configuration file is loaded on your switch, you do not need to do anything else.

Table 5-1 Prerequisites for Enabling Automatic Configuration

Device	Required Configuration
Access switch	Factory default (no configuration file)
Distribution switch	• IP helper address • Enable DHCP relay agent • IP routing (if used as default gateway)

Table 5-1 Prerequisites for Enabling Automatic Configuration (continued)

Device	Required Configuration
DHCP server	• IP address assignment • TFTP server IP address • Path to bootstrap configuration file on the TFTP server • Default gateway IP address
TFTP server	• A bootstrap configuration file that includes the CNS configuration commands that enable the switch to communicate with the Configuration Engine • The switch configured to use either the switch MAC address or the serial number (instead of the default hostname) to generate the ConfigID and EventID • The CNS event agent configured to push the configuration file to the switch
CNS Configuration Engine	One or more templates for each type of device, with the ConfigID of the device mapped to the template

Information About Configuring Cisco IOS Configuration Engine

Cisco Configuration Engine is network management software that acts as a configuration service for automating the deployment and management of network devices and services (see Figure 5-1). Each Cisco Configuration Engine service manages a group of Cisco devices (switches and routers) and the services that they deliver, storing their configurations and delivering them as needed.
Cisco Configuration Engine automates initial configurations and configuration updates by generating device-specific configuration changes, sending them to the device, executing the configuration change, and logging the results.

Cisco Configuration Engine supports standalone and server modes and has these CNS components:

- Configuration service (web server, file manager, and namespace mapping server)
- Event service (event gateway)
- Data service directory (data models and schema)

In standalone mode, Cisco Configuration Engine supports an embedded directory service. In this mode, no external directory or other data store is required. In server mode, Cisco Configuration Engine supports a user-defined external directory.

Figure 5-1 Configuration Engine Architectural Overview

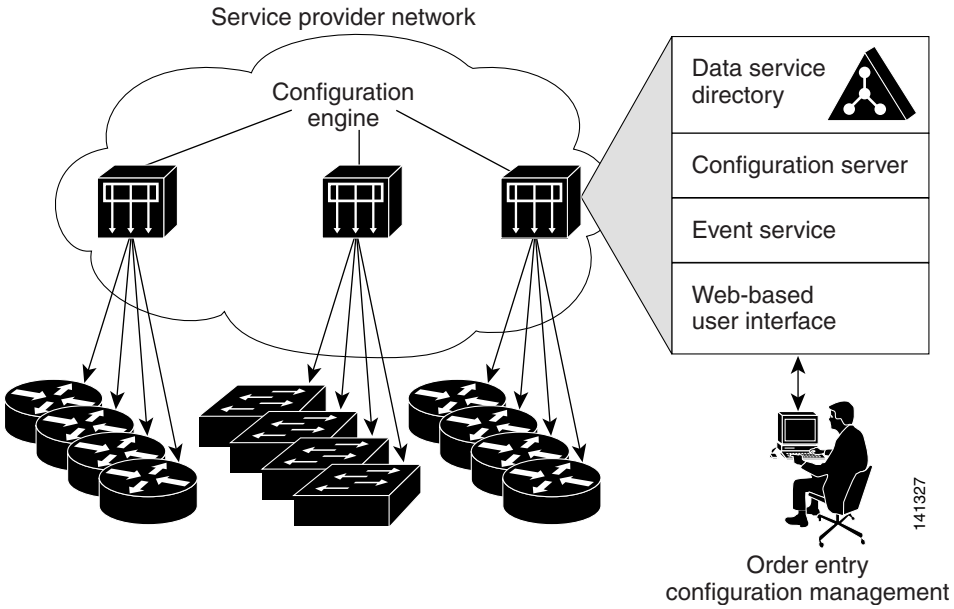

Configuration Service

Configuration Service is the core component of Cisco Configuration Engine. It consists of a configuration server that works with Cisco IOS CNS agents on the switch. Configuration Service delivers device and service configurations to the switch for initial configuration and mass reconfiguration by logical groups. Switches receive their initial configuration from the Configuration Service when they start up on the network for the first time.

Configuration Service uses CNS Event Service to send and receive configuration change events and to send success and failure notifications.

The configuration server is a web server that uses configuration templates and the device-specific configuration information stored in the embedded (standalone mode) or remote (server mode) directory.

Configuration templates are text files containing static configuration information in the form of CLI commands. In the templates, variables are specified using Lightweight Directory Access Protocol (LDAP) URLs that reference the device-specific configuration information stored in a directory.

The Cisco IOS agent can perform a syntax check on received configuration files and publish events to show the success or failure of the syntax check. The configuration agent can either apply configurations immediately or delay the application until receipt of a synchronization event from the configuration server.

Event Service

Cisco Configuration Engine uses Event Service for receipt and generation of configuration events. The event agent is on the switch and facilitates the communication between the switch and the event gateway on Configuration Engine.

Event Service is a highly capable publish-and-subscribe communication method. Event Service uses subject-based addressing to send messages to their destinations. Subject-based addressing conventions define a simple, uniform namespace for messages and their destinations.

NameSpace Mapper

Configuration Engine includes NameSpace Mapper (NSM), which provides a lookup service for managing logical groups of devices based on application, device or group ID, and event.

Cisco IOS devices recognize only event subject names that match those configured in Cisco IOS software; for example, cisco.cns.config.load. You can use the namespace mapping service to designate events by using any desired naming convention. When you have populated your data store with your subject names, NSM changes your event subject-name strings to those known by Cisco IOS.

For a subscriber, when given a unique device ID and event, the namespace mapping service returns a set of events to which to subscribe. Similarly, for a publisher, when given a unique group ID, device ID, and event, the mapping service returns a set of events on which to publish.

CNS IDs and Device Hostnames

Configuration Engine assumes that a unique identifier is associated with each configured switch. This unique identifier can take on multiple synonyms, where each synonym is unique within a particular namespace. The event service uses namespace content for subject-based addressing of messages.

Configuration Engine intersects two namespaces, one for the event bus and the other for the configuration server. Within the scope of the configuration server namespace, the term *ConfigID* is the unique identifier for a device. Within the scope of the event bus namespace, the term *DeviceID* is the CNS unique identifier for a device.

Because Configuration Engine uses both the event bus and the configuration server to provide configurations to devices, you must define both ConfigID and Device ID for each configured switch.

Within the scope of a single instance of the configuration server, no two configured switches can share the same value for ConfigID. Within the scope of a single instance of the event bus, no two configured switches can share the same value for DeviceID.

ConfigID

Each configured switch has a unique ConfigID, which serves as the key into the Configuration Engine directory for the corresponding set of switch CLI attributes. The ConfigID defined on the switch must match the ConfigID for the corresponding switch definition on Configuration Engine.

The ConfigID is fixed at startup time and cannot be changed until the device restarts, even if the switch hostname is reconfigured.

DeviceID

Each configured switch participating on the event bus has a unique DeviceID, which is analogous to the switch source address so that the switch can be targeted as a specific destination on the bus. All switches configured with the **cns config partial** global configuration command must access the event bus. Therefore, the DeviceID, as originated on the switch, must match the DeviceID of the corresponding switch definition in Configuration Engine.

The origin of the DeviceID is defined by the Cisco IOS hostname of the switch. However, the DeviceID variable and its usage reside within the event gateway adjacent to the switch.

The logical Cisco IOS termination point on the event bus is embedded in the event gateway, which in turn functions as a proxy on behalf of the switch. The event gateway represents the switch and its corresponding DeviceID to the event bus.

The switch declares its hostname to the event gateway immediately after the successful connection to the event gateway. The event gateway couples the DeviceID value to the Cisco IOS hostname each time this connection is established. The event gateway caches this DeviceID value for the duration of its connection to the switch.

Hostname and DeviceID Interaction

The DeviceID is fixed at the time of the connection to the event gateway and does not change even when the switch hostname is reconfigured.

When changing the switch hostname on the switch, the only way to refresh the DeviceID is to break the connection between the switch and the event gateway. Enter the **no cns event** global configuration command followed by the **cns event** global configuration command.

When the connection is reestablished, the switch sends its modified hostname to the event gateway. The event gateway redefines the DeviceID to the new value.

Using Hostname, DeviceID, and ConfigID

In standalone mode, when a hostname value is set for a switch, the configuration server uses the hostname as the DeviceID when an event is sent on hostname. If the hostname has not been set, the event is sent on the cn=<*value*> of the device.

In server mode, the hostname is not used. In this mode, the unique DeviceID attribute is always used for sending an event on the bus. If this attribute is not set, you cannot update the switch.

These and other associated attributes (tag value pairs) are set when you run **Setup** on Configuration Engine.

Cisco IOS Agents

The CNS event agent feature allows the switch to publish and subscribe to events on the event bus and works with the Cisco IOS agent.

Initial Configuration

When the switch first comes up, it attempts to get an IP address by broadcasting a DHCP request on the network. Assuming there is no DHCP server on the subnet, the distribution switch acts as a DHCP relay agent and forwards the request to the DHCP server. Upon receiving the request, the DHCP server assigns an IP address to the new switch and includes the TFTP server IP address, the path to the bootstrap configuration file, and the default gateway IP address in a unicast reply to the DHCP relay agent. The DHCP relay agent forwards the reply to the switch.

The switch automatically configures the assigned IP address on interface VLAN 1 (the default) and downloads the bootstrap configuration file from the TFTP server. Upon successful download of the bootstrap configuration file, the switch loads the file in its running configuration.

The Cisco IOS agents initiate communication with Configuration Engine by using the appropriate ConfigID and EventID. Configuration Engine maps the ConfigID to a template and downloads the full configuration file to the switch.

Figure 5-2 shows a sample network configuration for retrieving the initial bootstrap configuration file by using DHCP-based autoconfiguration.

Figure 5-2 Initial Configuration Overview

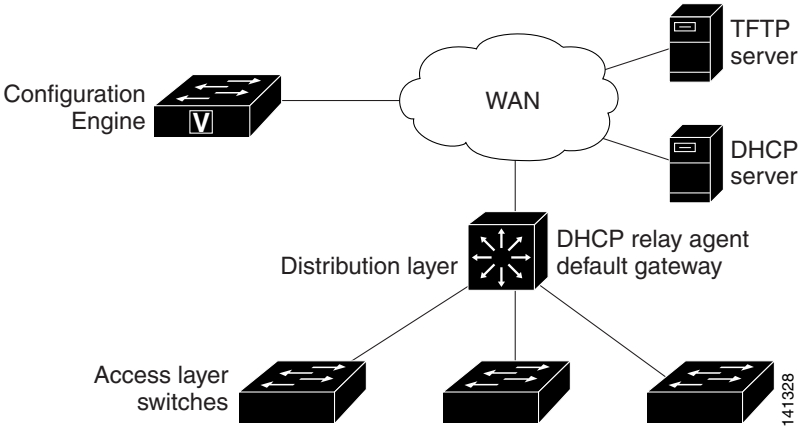

Incremental (Partial) Configuration

After the network is running, new services can be added by using the Cisco IOS agent. Incremental (partial) configurations can be sent to the switch. The actual configuration can be sent as an event payload by way of the event gateway (push operation) or as a signal event that triggers the switch to initiate a pull operation.

The switch can check the syntax of the configuration before applying it. If the syntax is correct, the switch applies the incremental configuration and publishes an event that signals success to the configuration server. If the switch does not apply the incremental configuration, it publishes an event showing an error status. When the switch has applied the incremental configuration, it can write it to NVRAM or wait until signaled to do so.

Synchronized Configuration

When the switch receives a configuration, it can defer application of the configuration upon receipt of a write-signal event. The write-signal event tells the switch not to save the updated configuration into its NVRAM. The switch uses the updated configuration as its running configuration. This ensures that the switch configuration is synchronized with other network activities before saving the configuration in NVRAM for use at the next reboot.

How to Configure Cisco IOS Configuration Engine

Configuring Cisco IOS Agents

CNS Event Agent and Cisco IOS CNS Agent embedded in the Cisco IOS software on the switch allows the switch to be connected and automatically configured. Both agents must be enabled and the CNS configuration can be **initial** or **partial**. The partial configuration allows you to use Configuration Engine to remotely send incremental configuration to the switch.

Enabling CNS Event Agent

Before You Begin

You must enable CNS Event Agent on the switch before you enable Cisco IOS CNS Agent.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**cns event** {*hostname* | *ip-address*} [*port-number*] [**backup**] [**failover-time** *seconds*] [**keepalive** *seconds retry-count*] [**reconnect** *time*] [**source** *ip-address*]	Enables the event agent, and enters the gateway parameters. • {*hostname* | *ip-address*}—Enters either the hostname or the IP address of the event gateway. • (Optional) *port number*—Enters the port number for the event gateway. The default port number is 11011. • (Optional) **backup**—Shows that this is the backup gateway. (If omitted, this is the primary gateway.) • (Optional) **failover-time** *seconds*—Enters how long the switch waits for the primary gateway route after the route to the backup gateway is established. • (Optional) **keepalive** *seconds*—Enters how often the switch sends keepalive messages. For *retry-count*, enters the number of unanswered keepalive messages that the switch sends before the connection is terminated. The default for each is 0. • (Optional) **reconnect** *time*—Enters the maximum time interval that the switch waits before trying to reconnect to the event gateway. • (Optional) **source** *ip-address*—Enters the source IP address of this device. **Note** Though visible in the command-line help string, the **encrypt** and the **clock-timeout** *time* keywords are not supported.
Step 3	**end**	Returns to privileged EXEC mode.
Step 4	**show cns event connections**	Verifies information about the event agent.

Enabling Cisco IOS CNS Agent and an Initial Configuration

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**cns template connect** *name*	Enters CNS template connect configuration mode, and specifies the name of the CNS connect template.
Step 3	**cli** *config-text*	Enters a command line for the CNS connect template. Repeat this step for each command line in the template.
Step 4		Repeat Steps 2 to 3 to configure another CNS connect template.
Step 5	**exit**	Returns to global configuration mode.
Step 6	**cns connect** *name* [**retries** *number*] [**retry-interval** *seconds*] [**sleep** *seconds*] [**timeout** *seconds*]	Enters CNS connect configuration mode, specifies the name of the CNS connect profile, and defines the profile parameters. The switch uses the CNS connect profile to connect to Configuration Engine. • (Optional) **retries** *number*—Enters the number of connection retries. The range is 1 to 30. The default is 3. • (Optional) **retry-interva**l *seconds*—Enters the interval between successive connection attempts to the Configuration Engine. The range is 1 to 40 seconds. The default is 10 seconds. • (Optional) **sleep** *seconds*—Enters the amount of time before which the first connection attempt occurs. The range is 0 to 250 seconds. The default is 0. • (Optional) **timeout** *seconds*—Enters the amount of time after which the connection attempts end. The range is 10 to 2000 seconds. The default is 120.
Step 7	**discover** {**controller** *controller-type* | **dlci** [**subinterface** *subinterface-number*] | **interface** [*interface-type*] | **line** *line-type*}	Specifies the interface parameters in the CNS connect profile. • **controller** *controller-type*—Enters the controller type. • **dlci**—Enters the active data-link connection identifiers (DLCIs). (Optional) **subinterface** *subinterface-number*—Specifies the point-to-point subinterface number that is used to search for active DLCIs. • **interface** [*interface-type*]—Enters the type of interface. • **line** *line-type*—Enters the line type.
Step 8	**template** *name* [... *name*]	Specifies the list of CNS connect templates in the CNS connect profile to be applied to the switch configuration. You can specify more than one template.
Step 9		Repeat Steps 7 to 8 to specify more interface parameters and CNS connect templates in the CNS connect profile.
Step 10	**exit**	Returns to global configuration mode.
Step 11	**hostname** *name*	Enters the hostname for the switch.

	Command	Purpose
Step 12	**ip route** *network-number*	(Optional) Establishes a static route to Configuration Engine whose IP address is *network-number*.
Step 13	**cns id** *interface num* {**dns-reverse** | **ipaddress** | **mac-address**} [**event**] [**image**] or **cns id** {**hardware-serial** | **hostname** | **string** *string* | **udi**} [**event**] [**image**]	(Optional) Sets the unique EventID or ConfigID used by the Configuration Engine. • *interface num*—Enters the type of interface for example, ethernet, group-async, loopback, or virtual-template. This setting specifies from which interface the IP or MAC address should be retrieved to define the unique ID. • **dns-reverse**—Retrieves the hostname and assigns it as the unique ID. • **ipaddress**—Uses the IP address. • **mac-address**—Uses the MAC address as the unique ID. • (Optional) **event**—Sets the ID to be the eventID value used to identify the switch. • (Optional) **image**—Sets the ID to be the imageID value used to identify the switch. **Note** If the **event** and **image** keywords are omitted, the imageID value is used to identify the switch. • **hardware-serial**—Sets the switch serial number as the unique ID. • **hostname** (the default)—Selects the switch hostname as the unique ID, uses an arbitrary text string **string** *string* as the unique ID and **udi** sets the unique device identifier (UDI) as the unique ID.

	Command	Purpose
Step 14	**cns config initial** {*hostname* | *ip-address*} [*port-number*] [**event**] [**no-persist**] [**page** *page*] [**source** *ip-address*] [**syntax-check**]	Enables the Cisco IOS agent and initiates an initial configuration. • {*hostname* | *ip-address*}—Enters the hostname or the IP address of the configuration server. • (Optional) *port-number*—Enters the port number of the configuration server. The default port number is 80. • (Optional) **event**—Enables configuration success, failure, or warning messages when the configuration is finished. • (Optional) **no-persist**—Suppresses the automatic writing to NVRAM of the configuration pulled as a result of entering the **cns config initial** global configuration command. If the **no-persist** keyword is not entered, using the **cns config initial** command causes the resultant configuration to be automatically written to NVRAM. • (Optional) **page** *page*—Enters the web page of the initial configuration. The default is /Config/config/asp. • (Optional) **source** *ip-address*—Enters the source IP address. • (Optional) **syntax-check**—Checks the syntax when this parameter is entered. **Note** Though visible in the command-line help string, the **encrypt**, **status** *url*, and **inventory** keywords are not supported.
Step 15	**end**	Returns to privileged EXEC mode.

Enabling a Partial Configuration

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**cns config partial** {*ip-address* | *hostname*} [*port-number*] [**source** *ip-address*]	Enables the configuration agent, and initiates a partial configuration. • {*ip-address* | *hostname*}—Enters the IP address or the hostname of the configuration server. • (Optional) *port-number*—Enters the port number of the configuration server. The default port number is 80. • (Optional) **source** *ip-address*—Enters the source IP address. **Note** Though visible in the command-line help string, the **encrypt** keyword is not supported.
Step 3	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining Cisco IOS Configuration Engine

Command	Purpose
show cns config connections	Displays the status of the CNS Cisco IOS agent connections.
show cns config outstanding	Displays information about incremental (partial) CNS configurations that have started but are not yet completed.
show cns config stats	Displays statistics about the Cisco IOS agent.
show cns event connections	Displays the status of the CNS event agent connections.
show cns event stats	Displays statistics about the CNS event agent.
show cns event subject	Displays a list of event agent subjects that are subscribed to by applications.

Configuration Examples for Cisco IOS Configuration Engine

Enabling the CNS Event Agent: Example

This example shows how to enable the CNS event agent, set the IP address gateway to 10.180.1.27, set 120 seconds as the keepalive interval, and set 10 as the retry count.

```
Switch(config)# cns event 10.180.1.27 keepalive 120 10
```

Configuring an Initial CNS Configuration: Examples

This example shows how to configure an initial configuration on a remote switch when the switch configuration is unknown (the CNS Zero Touch feature).

```
Switch(config)# cns template connect template-dhcp
Switch(config-tmpl-conn)# cli ip address dhcp
Switch(config-tmpl-conn)# exit
Switch(config)# cns template connect ip-route
Switch(config-tmpl-conn)# cli ip route 0.0.0.0 0.0.0.0 ${next-hop}
Switch(config-tmpl-conn)# exit
Switch(config)# cns connect dhcp
Switch(config-cns-conn)# discover interface gigabitethernet
Switch(config-cns-conn)# template template-dhcp
Switch(config-cns-conn)# template ip-route
Switch(config-cns-conn)# exit
Switch(config)# hostname RemoteSwitch
RemoteSwitch(config)# cns config initial 10.1.1.1 no-persist
```

This example shows how to configure an initial configuration on a remote switch when the switch IP address is known. The Configuration Engine IP address is 172.28.129.22.

```
Switch(config)# cns template connect template-dhcp
Switch(config-tmpl-conn)# cli ip address dhcp
Switch(config-tmpl-conn)# exit
Switch(config)# cns template connect ip-route
Switch(config-tmpl-conn)# cli ip route 0.0.0.0 0.0.0.0 ${next-hop}
Switch(config-tmpl-conn)# exit
Switch(config)# cns connect dhcp
```

```
Switch(config-cns-conn)# discover interface gigabitethernet
Switch(config-cns-conn)# template template-dhcp
Switch(config-cns-conn)# template ip-route
Switch(config-cns-conn)# exit
Switch(config)# hostname RemoteSwitch
RemoteSwitch(config)# ip route 172.28.129.22 255.255.255.255 11.11.11.1
RemoteSwitch(config)# cns id ethernet 0 ipaddress
RemoteSwitch(config)# cns config initial 172.28.129.22 no-persist
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
Network management commands	*Cisco IOS Network Management Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 6

Configuring Switch Clusters

This chapter provides the concepts and procedures to create and manage switch clusters on your switch. You can create and manage switch clusters by using Cisco Network Assistant (CNA), the command-line interface (CLI), or SNMP. For complete procedures, see the online help for CNA. For the CLI cluster commands, see the switch command reference.

This chapter focuses on Cisco IE 2000 switch clusters. It also includes guidelines and limitations for clusters mixed with other cluster-capable Catalyst switches, but it does not provide complete descriptions of the cluster features for switches in the cluster. For complete cluster information for a specific Catalyst platform, refer to the software configuration guide for that switch.

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for Configuring Switch Clusters

- Static routing and routed ports is supported only when the Cisco IE 2000 switch is running the LAN Base image.

Cluster Command Switch Characteristics

A cluster command switch must meet these requirements:

- Is running Cisco IOS Release 15.0(1)EY or later.
- Has an IP address.
- Has Cisco Discovery Protocol (CDP) version 2 enabled (the default).
- Is not a command or cluster member switch of another cluster.
- Is connected to the standby cluster command switches through the management VLAN and to the cluster member switches through a common VLAN.

Standby Cluster Command Switch Characteristics

A standby cluster command switch must meet these requirements:

- Is running Cisco IOS 15.0(1)EY or later.
- Has an IP address.
- Has CDP version 2 enabled.
- Is connected to the command switch and to other standby command switches through its management VLAN.
- Is connected to all other cluster member switches (except the cluster command and standby command switches) through a common VLAN.
- Is redundantly connected to the cluster so that connectivity to cluster member switches is maintained.
- Is not a command or member switch of another cluster.

Candidate Switch and Cluster Member Switch Characteristics

Candidate switches are cluster-capable switches that have not yet been added to a cluster. Cluster member switches are switches that have actually been added to a switch cluster. Although not required, a candidate or cluster member switch can have its own IP address and password (for related considerations, see the "IP Addresses" section on page 6-11 and "Passwords" section on page 6-12).

To join a cluster, a candidate switch must meet these requirements:

- Is running cluster-capable software.
- Has CDP version 2 enabled.
- Is not a command or cluster member switch of another cluster.
- If a cluster standby group exists, the switch is connected to every standby cluster command switch through at least one common VLAN. The VLAN to each standby cluster command switch can be different.
- Is connected to the cluster command switch through at least one common VLAN.

Note Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL candidate and cluster member switches must be connected through their management VLAN to the cluster command switch and standby cluster command switches. For complete information about these switches in a switch-cluster environment, refer to the software configuration guide for that specific switch.

This requirement does not apply if you have a Catalyst 2970, Catalyst 3550, Catalyst 3560, or Catalyst 3750 cluster command switch. Candidate and cluster member switches can connect through any VLAN in common with the cluster command switch.

Restrictions for Configuring Switch Clusters

We do not recommend using the **ip http access-class** global configuration command to limit access to specific hosts or networks. Access should be controlled through the cluster command switch or by applying access control lists (ACLs) on interfaces that are configured with IP address. For more information on ACLs, see Chapter 37, "Configuring Network Security with ACLs."

Information About Configuring Switch Clusters

A *switch cluster* is a set of up to 16 connected, cluster-capable Catalyst switches that are managed as a single entity. The switches in the cluster use the switch clustering technology so that you can configure and troubleshoot a group of different Catalyst desktop switch platforms through a single IP address.

In a switch cluster, one switch must be the *cluster command switch* and up to 15 other switches can be *cluster member switches*. The total number of switches in a cluster cannot exceed 16 switches. The cluster command switch is the single point of access used to configure, manage, and monitor the cluster member switches. Cluster members can belong to only one cluster at a time.

Benefits of Clustering Switches

- Management of switches regardless of their interconnection media and their physical locations. The switches can be in the same location, or they can be distributed across a Layer 2 or Layer 3 (if your cluster is using a Catalyst 3550, Catalyst 3560, or Catalyst 3750 switch as a Layer 3 router between the Layer 2 switches in the cluster) network.

  Cluster members are connected to the cluster command switch according to the connectivity guidelines described in the "Automatic Discovery of Cluster Candidates and Members" section on page 6-5. This section includes management VLAN considerations for the Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL switches. For complete information about these switches in a switch-cluster environment, refer to the software configuration guide for that specific switch.
- Command-switch redundancy if a cluster command switch fails. One or more switches can be designated as *standby cluster command switches* to avoid loss of contact with cluster members. A *cluster standby group* is a group of standby cluster command switches.
- Management of a variety of switches through a single IP address. This preserves IP addresses, especially if you have a limited number of them. All communication with the switch cluster is through the cluster command switch IP address.

Eligible Cluster Switches

Table 6-1 lists the switches eligible for switch clustering, including which ones can be cluster command switches and which ones can only be cluster member switches, and the required software versions.

Table 6-1 Switch Software and Cluster Capability

Switch	Cisco IOS Release	Cluster Capability
IE 2000 switch	15.0(1)EY or later	Member or command switch
IE 3010 switch	12.2(53)EZ or later	Member or command switch

Table 6-1 Switch Software and Cluster Capability (continued)

Switch	Cisco IOS Release	Cluster Capability
IE 3000 switch	12.2(40)EX or later	Member or command switch
Catalyst 3750-X or Catalyst 3560-X	12.2(53)SE2 or later	Member or command switch
Catalyst 3750-E or Catalyst 3560-E	12.2(35)SE2 or later	Member or command switch
Catalyst 3750	12.1(11)AX or later	Member or command switch
Catalyst 3560	12.1(19)EA1b or later	Member or command switch
Catalyst 3550	12.1(4)EA1 or later	Member or command switch
Catalyst 2975	12.2(46)EX or later	Member or command switch
Catalyst 2970	12.1(11)AX or later	Member or command switch
Catalyst 2960-S	12.2(53)SE or later	Member or command switch
Catalyst 2960	12.2(25)FX or later	Member or command switch
Catalyst 2955	12.1(12c)EA1 or later	Member or command switch
Catalyst 2950	12.0(5.2)WC(1) or later	Member or command switch
Catalyst 2950 LRE	12.1(11)JY or later	Member or command switch
Catalyst 2940	12.1(13)AY or later	Member or command switch
Catalyst 3500 XL	12.0(5.1)XU or later	Member or command switch
Catalyst 2900 XL (8-MB switches)	12.0(5.1)XU or later	Member or command switch
Catalyst 2900 XL (4-MB switches)	11.2(8.5)SA6 (recommended)	Member switch only
Catalyst 1900 and 2820	9.00(-A or -EN) or later	Member switch only

How to Plan for Switch Clustering

Anticipating conflicts and compatibility issues is a high priority when you manage several switches through a cluster. This section describes the guidelines, requirements, and caveats that you should understand before you create the cluster:

- Automatic Discovery of Cluster Candidates and Members, page 6-5
- IP Addresses, page 6-11
- Hostnames, page 6-11
- Passwords, page 6-12
- SNMP Community Strings, page 6-12
- TACACS+ and RADIUS, page 6-12
- LRE Profiles, page 6-13

Refer to the release notes for the list of Catalyst switches eligible for switch clustering, including which ones can be cluster command switches and which ones can only be cluster member switches, and for the required software versions and browser and Java plug-in configurations.

Automatic Discovery of Cluster Candidates and Members

The cluster command switch uses Cisco Discovery Protocol (CDP) to discover cluster member switches, candidate switches, neighboring switch clusters, and edge devices across multiple VLANs and in star or cascaded topologies.

Note Do not disable CDP on the cluster command switch, on cluster members, or on any cluster-capable switches that you might want a cluster command switch to discover. For more information about CDP, see Chapter 32, "Configuring CDP."

Following these connectivity guidelines ensures automatic discovery of the switch cluster, cluster candidates, connected switch clusters, and neighboring edge devices:

- Discovery Through CDP Hops, page 6-5
- Discovery Through Non-CDP-Capable and Noncluster-Capable Devices, page 6-7
- Discovery Through Different VLANs, page 6-7
- Discovery Through Different Management VLANs, page 6-8
- Discovery Through Routed Ports, page 6-9
- Discovery of Newly Installed Switches, page 6-10

Discovery Through CDP Hops

By using CDP, a cluster command switch can discover switches up to seven CDP hops away (the default is three hops) from the edge of the cluster. The edge of the cluster is where the last cluster member switches are connected to the cluster and to candidate switches. For example, cluster member switches 9 and 10 in Figure 6-1 are at the edge of the cluster.

In Figure 6-1, the cluster command switch has ports assigned to VLANs 16 and 62. The CDP hop count is three. The cluster command switch discovers switches 11, 12, 13, and 14 because they are within three hops from the edge of the cluster. It does not discover switch 15 because it is four hops from the edge of the cluster.

Figure 6-1 Discovery Through CDP Hops

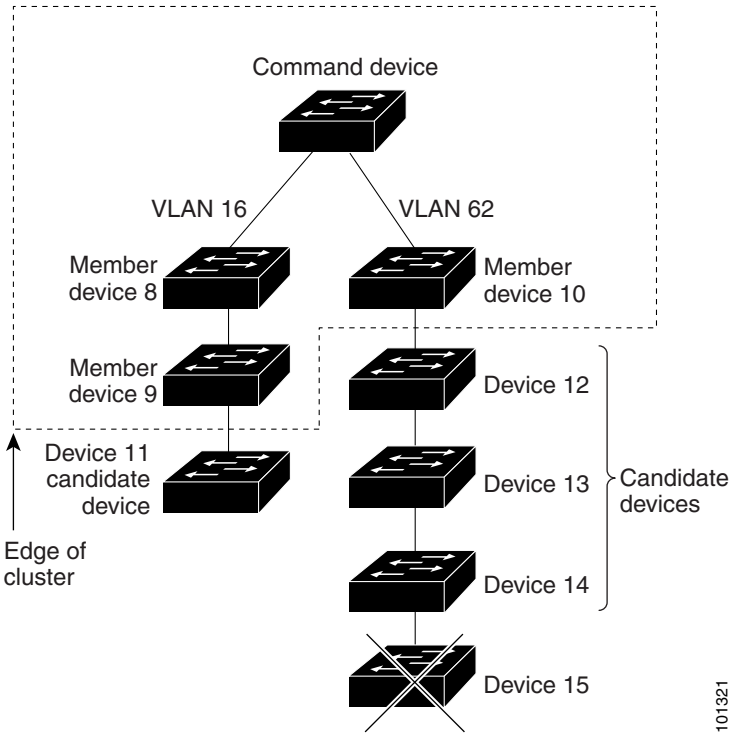

Discovery Through Non-CDP-Capable and Noncluster-Capable Devices

If a cluster command switch is connected to a *non-CDP-capable third-party hub* (such as a non-Cisco hub), it can discover cluster-enabled devices connected to that third-party hub. However, if the cluster command switch is connected to a *noncluster-capable Cisco device*, it cannot discover a cluster-enabled device connected beyond the noncluster-capable Cisco device.

Figure 6-2 shows that the cluster command switch discovers the switch that is connected to a third-party hub. However, the cluster command switch does not discover the switch that is connected to a Catalyst 5000 switch.

Figure 6-2 Discovery Through Non-CDP-Capable and Noncluster-Capable Devices

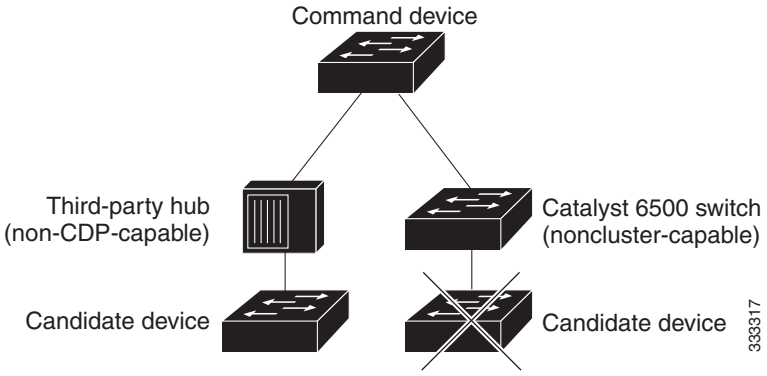

Discovery Through Different VLANs

If the cluster command switch is a Catalyst 2970, Catalyst 3550, Catalyst 3560, or Catalyst 3750 switch, the cluster can have cluster member switches in different VLANs. As cluster member switches, they must be connected through at least one VLAN in common with the cluster command switch. The cluster command switch in Figure 6-3 has ports assigned to VLANs 9, 16, and 62 and therefore discovers the switches in those VLANs. It does not discover the switch in VLAN 50. It also does not discover the switch in VLAN 16 in the first column because the cluster command switch has no VLAN connectivity to it.

Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL cluster member switches must be connected to the cluster command switch through their management VLAN. For information about discovery through management VLANs, see the "Discovery Through Different Management VLANs" section on page 6-8. For more information about VLANs, see Chapter 17, "Configuring VLANs."

Figure 6-3 Discovery Through Different VLANs

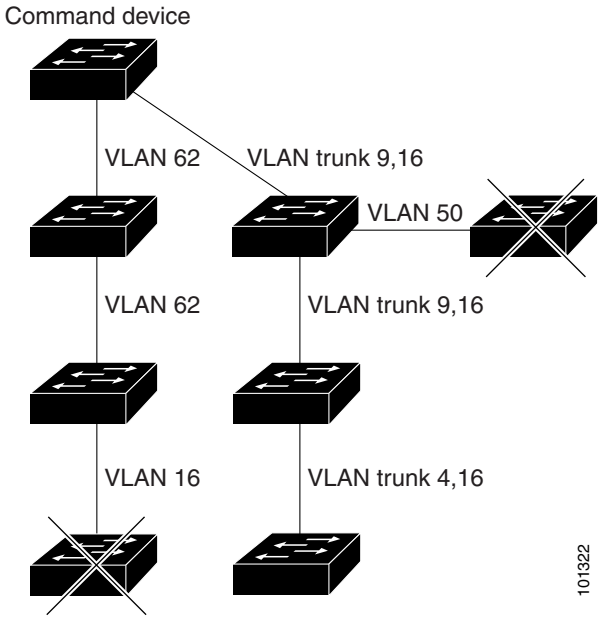

Discovery Through Different Management VLANs

Catalyst 2970, Catalyst 3550, Catalyst 3560, or Catalyst 3750 cluster command switches can discover and manage cluster member switches in different VLANs and different management VLANs. As cluster member switches, they must be connected through at least one VLAN in common with the cluster command switch. They do not need to be connected to the cluster command switch through their management VLAN. The default management VLAN is VLAN 1.

Note If the switch cluster has a Catalyst 3750 or 2975 switch or has a switch stack, that switch or switch stack must be the cluster command switch.

The cluster command switch and standby command switch in Figure 6-5 (assuming they are Catalyst 2960, Catalyst 2970, Catalyst 2975, Catalyst 3550, Catalyst 3560, or Catalyst 3750 cluster command switches) have ports assigned to VLANs 9, 16, and 62. The management VLAN on the cluster command switch is VLAN 9. Each cluster command switch discovers the switches in the different management VLANs except these:

- Switches 7 and 10 (switches in management VLAN 4) because they are not connected through a common VLAN (meaning VLANs 62 and 9) with the cluster command switch
- Switch 9 because automatic discovery does not extend beyond a noncandidate device, which is switch 7

Discovery Through Routed Ports

Note The LAN Base image supports static routing and RIP.

If the cluster command switch has a routed port (RP) configured, it discovers only candidate and cluster member switches in the *same* VLAN as the routed port.

The Layer 3 cluster command switch in Figure 6-4 can discover the switches in VLANs 9 and 62 but not the switch in VLAN 4. If the routed port path between the cluster command switch and cluster member switch 7 is lost, connectivity with cluster member switch 7 is maintained because of the redundant path through VLAN 9.

Figure 6-4 Discovery Through Routed Ports

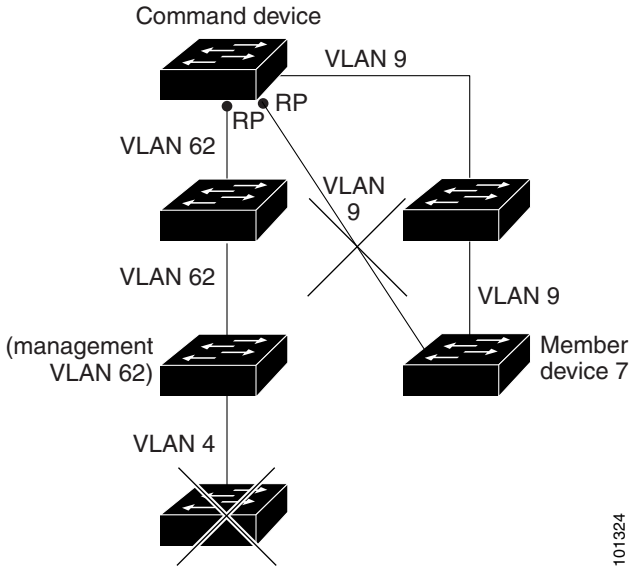

Figure 6-5 Discovery Through Different Management VLANs with a Layer 3 Cluster Command Switch

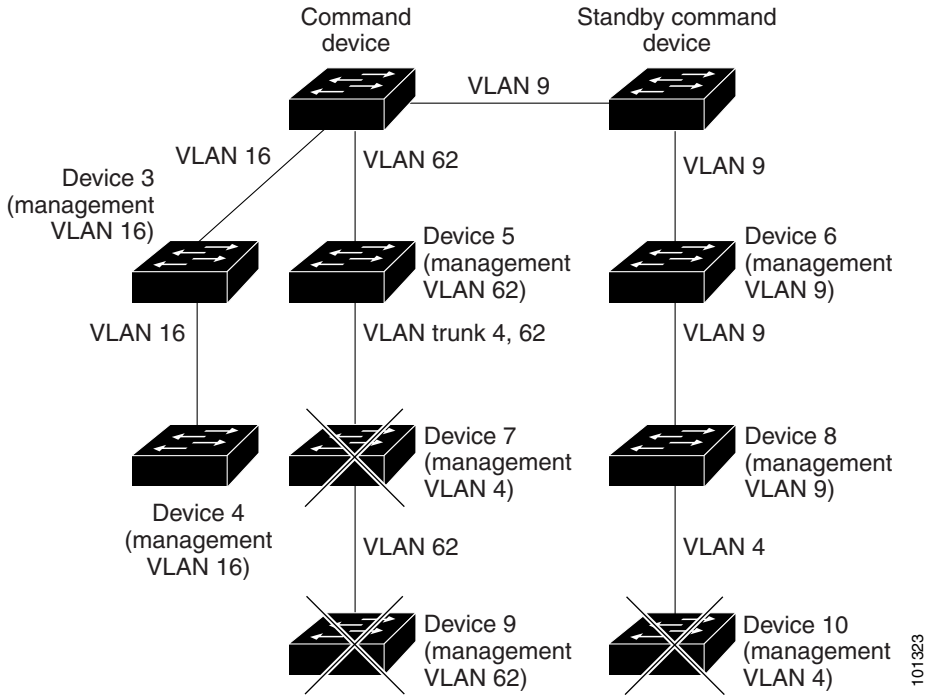

Discovery of Newly Installed Switches

To join a cluster, the new, out-of-the-box switch must be connected to the cluster through one of its access ports. An access port (AP) carries the traffic of and belongs to only one VLAN. By default, the new switch and its access ports are assigned to VLAN 1.

When the new switch joins a cluster, its default VLAN changes to the VLAN of the immediately upstream neighbor. The new switch also configures its access port to belong to the VLAN of the immediately upstream neighbor.

The cluster command switch in Figure 6-6 belongs to VLANs 9 and 16. When new cluster-capable switches join the cluster:

- One cluster-capable switch and its access port are assigned to VLAN 9.
- The other cluster-capable switch and its access port are assigned to management VLAN 16.

Figure 6-6 Discovery of Newly Installed Switches

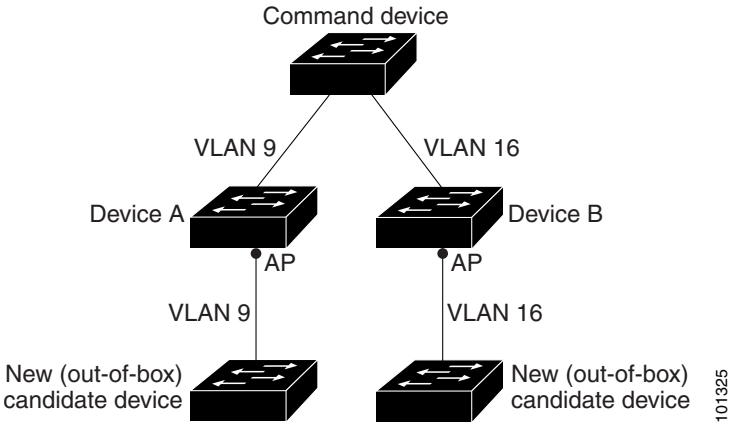

IP Addresses

You must assign IP information to a cluster command switch. You can assign more than one IP address to the cluster command switch, and you can access the cluster through any of the command-switch IP addresses. If you configure a cluster standby group, you must use the standby-group virtual IP address to manage the cluster from the active cluster command switch. Using the virtual IP address ensures that you retain connectivity to the cluster if the active cluster command switch fails and that a standby cluster command switch becomes the active cluster command switch.

If the active cluster command switch fails and the standby cluster command switch takes over, you must either use the standby-group virtual IP address or any of the IP addresses available on the new active cluster command switch to access the cluster.

You can assign an IP address to a cluster-capable switch, but it is not necessary. A cluster member switch is managed and communicates with other cluster member switches through the command-switch IP address. If the cluster member switch leaves the cluster and it does not have its own IP address, you must assign an IP address to manage it as a standalone switch.

For more information about IP addresses, see Chapter 4, “Performing Switch Setup Configuration.”

Hostnames

You do not need to assign a hostname to either a cluster command switch or an eligible cluster member. However, a hostname assigned to the cluster command switch can help to identify the switch cluster. The default hostname for the switch is *Switch*.

If a switch joins a cluster and it does not have a hostname, the cluster command switch appends a unique member number to its own hostname and assigns it sequentially as each switch joins the cluster. The number means the order in which the switch was added to the cluster. For example, a cluster command switch named *eng-cluster* could name the fifth cluster member *eng-cluster-5*.

If a switch has a hostname, it retains that name when it joins a cluster and when it leaves the cluster.

If a switch received its hostname from the cluster command switch, was removed from a cluster, was then added to a new cluster, and kept the same member number (such as *5*), the switch overwrites the old hostname (such as *eng-cluster-5*) with the hostname of the cluster command switch in the new cluster (such as *mkg-cluster-5*). If the switch member number changes in the new cluster (such as *3*), the switch retains the previous name (*eng-cluster-5*).

Passwords

You do not need to assign passwords to an individual switch if it will be a cluster member. When a switch joins a cluster, it inherits the command-switch password and retains it when it leaves the cluster. If no command-switch password is configured, the cluster member switch inherits a null password. Cluster member switches only inherit the command-switch password.

If you change the member-switch password to be different from the command-switch password and save the change, the switch is not manageable by the cluster command switch until you change the member-switch password to match the command-switch password. Rebooting the member switch does not revert the password back to the command-switch password. We recommend that you do not change the member-switch password after it joins a cluster.

For more information about passwords, see the "Prevention for Unauthorized Switch Access" section on page 12-2.

For password considerations specific to the Catalyst 1900 and Catalyst 2820 switches, refer to the installation and configuration guides for those switches.

SNMP Community Strings

A cluster member switch inherits the command-switch first read-only (RO) and read-write (RW) community strings with *@esN* appended to the community strings:

- *command-switch-readonly-community-string@esN*, where *N* is the member-switch number.
- *command-switch-readwrite-community-string@esN*, where *N* is the member-switch number.

If the cluster command switch has multiple read-only or read-write community strings, only the first read-only and read-write strings are propagated to the cluster member switch.

The switches support an unlimited number of community strings and string lengths. For more information about SNMP and community strings, see Chapter 36, "Configuring SNMP."

For SNMP considerations specific to the Catalyst 1900 and Catalyst 2820 switches, refer to the installation and configuration guides specific to those switches.

TACACS+ and RADIUS

If TACACS+ is configured on a cluster member, it must be configured on all cluster members. Similarly, if RADIUS is configured on a cluster member, it must be configured on all cluster members.The same switch cluster cannot have some members configured with TACACS+ and other members configured with RADIUS.

For more information about TACACS+, see the "Configuring TACACS+" section on page 12-30. For more information about RADIUS, see the "Configuring Radius Server Communication" section on page 12-33.

LRE Profiles

A configuration conflict occurs if a switch cluster has Long-Reach Ethernet (LRE) switches that use both private and public profiles. If one LRE switch in a cluster is assigned a public profile, all LRE switches in that cluster must have that same public profile. Before you add an LRE switch to a cluster, make sure that you assign it the same public profile used by other LRE switches in the cluster.

A cluster can have a mix of LRE switches that use different private profiles.

Managing Switch Clusters

Using the CLI to Manage Switch Clusters

You can configure cluster member switches from the CLI by first logging into the cluster command switch. Enter the **rcommand** user EXEC command and the cluster member switch number to start a Telnet session (through a console or Telnet connection) and to access the cluster member switch CLI. The command mode changes, and the Cisco IOS commands operate as usual. Enter the **exit** privileged EXEC command on the cluster member switch to return to the command-switch CLI.

This example shows how to log into member-switch 3 from the command-switch CLI:

```
switch# rcommand 3
```

If you do not know the member-switch number, enter the **show cluster members** privileged EXEC command on the cluster command switch. For more information about the **rcommand** command and all other cluster commands, refer to the switch command reference.

The Telnet session accesses the member-switch CLI at the same privilege level as on the cluster command switch. The Cisco IOS commands then operate as usual. For instructions on configuring the switch for a Telnet session, see the "Disabling Password Recovery" section on page 12-27.

Catalyst 1900 and Catalyst 2820 CLI Considerations

If your switch cluster has Catalyst 1900 and Catalyst 2820 switches running standard edition software, the Telnet session accesses the management console (a menu-driven interface) if the cluster command switch is at privilege level 15. If the cluster command switch is at privilege level 1 to 14, you are prompted for the password to access the menu console.

Command-switch privilege levels map to the Catalyst 1900 and Catalyst 2820 cluster member switches running standard and Enterprise Edition Software as follows:

- If the command-switch privilege level is 1 to 14, the cluster member switch is accessed at privilege level 1.
- If the command-switch privilege level is 15, the cluster member switch is accessed at privilege level 15.

Note The Catalyst 1900 and Catalyst 2820 CLI is available only on switches running Enterprise Edition Software.

For more information about the Catalyst 1900 and Catalyst 2820 switches, refer to the installation and configuration guides for those switches.

Using SNMP to Manage Switch Clusters

When you first power on the switch, SNMP is enabled if you enter the IP information by using the setup program and accept its proposed configuration. If you did not use the setup program to enter the IP information and SNMP was not enabled, you can enable it as described in the Chapter 36, "Configuring SNMP." On Catalyst 1900 and Catalyst 2820 switches, SNMP is enabled by default.

When you create a cluster, the cluster command switch manages the exchange of messages between cluster member switches and an SNMP application. The cluster software on the cluster command switch appends the cluster member switch number (*@esN*, where *N* is the switch number) to the first configured read-write and read-only community strings on the cluster command switch and propagates them to the cluster member switch. The cluster command switch uses this community string to control the forwarding of gets, sets, and get-next messages between the SNMP management station and the cluster member switches.

Note When a cluster standby group is configured, the cluster command switch can change without your knowledge. Use the first read-write and read-only community strings to communicate with the cluster command switch if there is a cluster standby group configured for the cluster.

If the cluster member switch does not have an IP address, the cluster command switch redirects traps from the cluster member switch to the management station, as shown in Figure 6-7. If a cluster member switch has its own IP address and community strings, the cluster member switch can send traps directly to the management station, without going through the cluster command switch.

If a cluster member switch has its own IP address and community strings, they can be used in addition to the access provided by the cluster command switch. For more information about SNMP and community strings, see Chapter 36, "Configuring SNMP."

Figure 6-7 SNMP Management for a Cluster

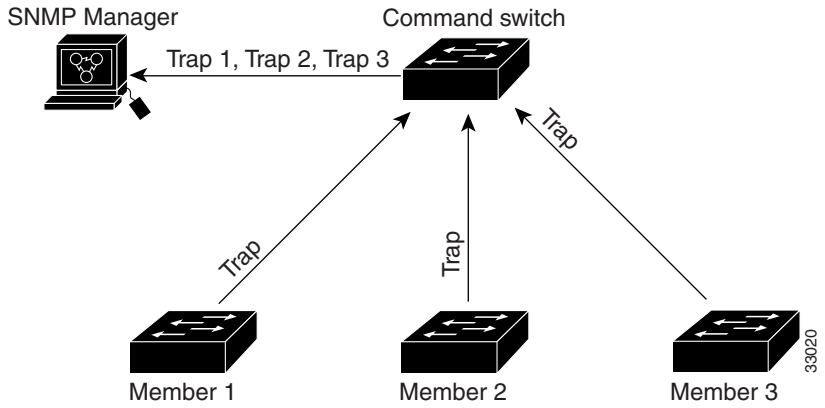

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 7

Performing Switch Administration

This chapter describes how to perform one-time operations to administer your switch.

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Information About Performing Switch Administration

System Time and Date Management

You can manage the system time and date on your switch using automatic configuration, such as the Network Time Protocol (NTP), or manual configuration methods.

System Clock

The basis of time service is the system clock. This clock runs from the moment the system starts up and keeps track of the date and time.

The system clock can then be set from these sources:

- NTP
- Manual configuration

The system clock can provide time to these services:

- User **show** commands
- Logging and debugging messages

The system clock keeps track of time internally based on Universal Time Coordinated (UTC), also known as Greenwich Mean Time (GMT). You can configure information about the local time zone and summer time (daylight saving time) so that the time appears correctly for the local time zone.

The system clock keeps track of whether the time is *authoritative* or not (that is, whether it has been set by a time source considered to be authoritative). If it is not authoritative, the time is available only for display purposes and is not redistributed. For configuration information, see the "Configuring Time and Date Manually" section on page 7-9.

Network Time Protocol

NTP is designed to time-synchronize a network of devices. NTP runs over User Datagram Protocol (UDP), which runs over IP. NTP is documented in RFC 1305.

An NTP network usually gets its time from an authoritative time source, such as a radio clock or an atomic clock attached to a time server. NTP then distributes this time across the network. NTP is extremely efficient; no more than one packet per minute is necessary to synchronize two devices to within a millisecond of one another.

NTP uses the concept of a *stratum* to describe how many NTP hops away a device is from an authoritative time source. A stratum 1 time server has a radio or atomic clock directly attached, a stratum 2 time server receives its time through NTP from a stratum 1 time server, and so on. A device running NTP automatically chooses as its time source the device with the lowest stratum number with which it communicates through NTP. This strategy effectively builds a self-organizing tree of NTP speakers.

NTP avoids synchronizing to a device whose time might not be accurate by never synchronizing to a device that is not synchronized. NTP also compares the time reported by several devices and does not synchronize to a device whose time is significantly different than the others, even if its stratum is lower.

The communications between devices running NTP (known as *associations*) are usually statically configured; each device is given the IP address of all devices with which it should form associations. Accurate timekeeping is possible by exchanging NTP messages between each pair of devices with an association. However, in a LAN environment, NTP can be configured to use IP broadcast messages instead. This alternative reduces configuration complexity because each device can simply be configured to send or receive broadcast messages. However, in that case, information flow is one-way only.

The time kept on a device is a critical resource; you should use the security features of NTP to avoid the accidental or malicious setting of an incorrect time. Two mechanisms are available: an access list-based restriction scheme and an encrypted authentication mechanism.

Cisco's implementation of NTP does not support stratum 1 service; it is not possible to connect to a radio or atomic clock. We recommend that the time service for your network be derived from the public NTP servers available on the IP Internet.

Figure 7-1 shows a typical network example using NTP. Switch A is the NTP master, with Switches B, C, and D configured in NTP server mode, in server association with Switch A. Switch E is configured as an NTP peer to the upstream and downstream switches, Switch B and Switch F.

Figure 7-1 Typical NTP Network Configuration

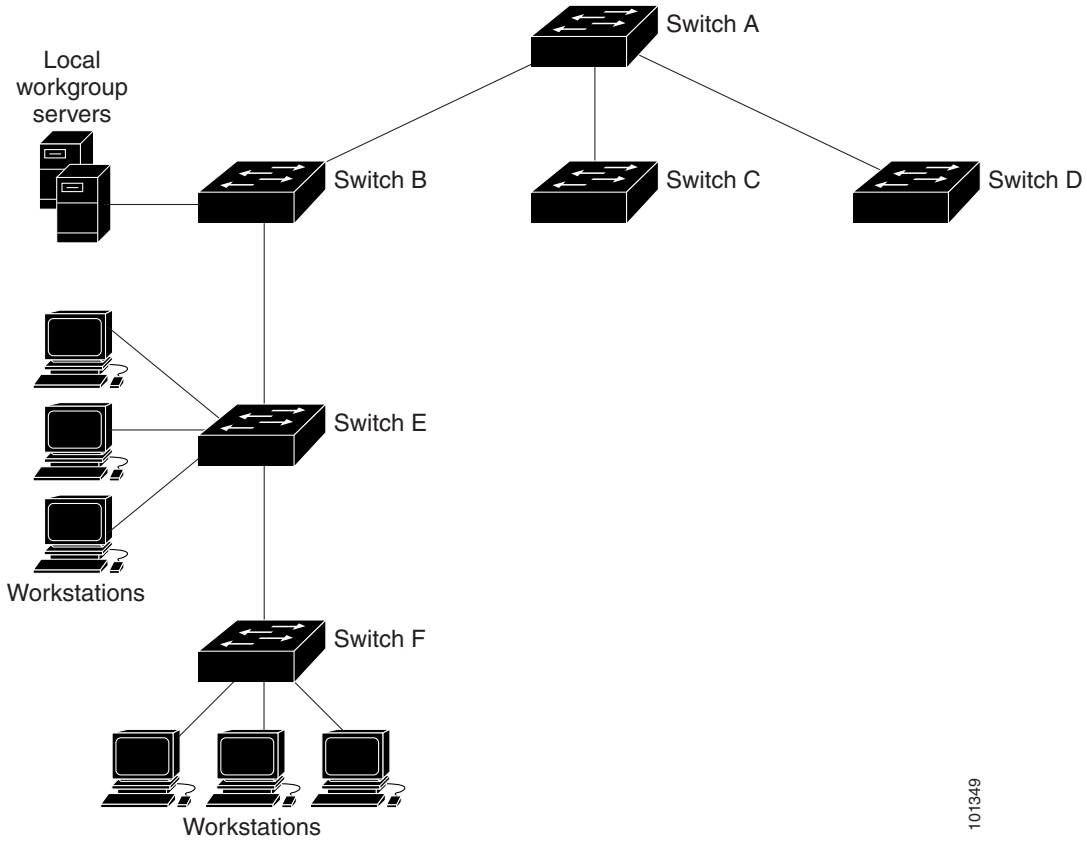

If the network is isolated from the Internet, Cisco's implementation of NTP allows a device to act as if it is synchronized through NTP, when in fact it has learned the time by using other means. Other devices then synchronize to that device through NTP.

When multiple sources of time are available, NTP is always considered to be more authoritative. NTP time overrides the time set by any other method.

Several manufacturers include NTP software for their host systems, and a publicly available version for systems running UNIX and its various derivatives is also available. This software allows host systems to be time-synchronized as well.

NTP Version 4

NTP version 4 is implemented on the switch. NTPv4 is an extension of NTP version 3. NTPv4 supports both IPv4 and IPv6 and is backward-compatible with NTPv3.

NTPv4 provides these capabilities:

- Support for IPv6.

- Improved security compared to NTPv3. The NTPv4 protocol provides a security framework based on public key cryptography and standard X509 certificates.
- Automatic calculation of the time-distribution hierarchy for a network. Using specific multicast groups, NTPv4 automatically configures the hierarchy of the servers to achieve the best time accuracy for the lowest bandwidth cost. This feature leverages site-local IPv6 multicast addresses.

For details about configuring NTPv4, see the *Cisco IOS IPv6 Configuration Guide* on Cisco.com.

DNS

The DNS protocol controls the Domain Name System (DNS), a distributed database with which you can map hostnames to IP addresses. When you configure DNS on your switch, you can substitute the hostname for the IP address with all IP commands, such as **ping**, **telnet**, **connect**, and related Telnet support operations.

IP defines a hierarchical naming scheme that allows a device to be identified by its location or domain. Domain names are pieced together with periods (.) as the delimiting characters. For example, Cisco Systems is a commercial organization that IP identifies by a *com* domain name, so its domain name is *cisco.com*. A specific device in this domain, for example, the File Transfer Protocol (FTP) system is identified as *ftp.cisco.com*.

To keep track of domain names, IP has defined the concept of a domain name server, which holds a cache (or database) of names mapped to IP addresses. To map domain names to IP addresses, you must first identify the hostnames, specify the name server that is present on your network, and enable the DNS.

Default DNS Configuration

Table 7-1 shows the default DNS configuration.

Table 7-1 Default DNS Configuration

Feature	Default Setting
DNS enable state	Enabled.
DNS default domain name	None configured.
DNS servers	No name server addresses are configured.

Login Banners

You can configure a message-of-the-day (MOTD) and a login banner. The MOTD banner displays on all connected terminals at login and is useful for sending messages that affect all network users (such as impending system shutdowns).

The login banner also displays on all connected terminals. It appears after the MOTD banner and before the login prompts.

The MOTD and login banners are not configured.

System Name and Prompt

You configure the system name on the switch to identify it. By default, the system name and prompt are *Switch*.

If you have not configured a system prompt, the first 20 characters of the system name are used as the system prompt. A greater-than symbol [>] is appended. The prompt is updated whenever the system name changes.

MAC Address Table

The MAC address table contains address information that the switch uses to forward traffic between ports. All MAC addresses in the address table are associated with one or more ports. The address table includes these types of addresses:

- Dynamic address—A source MAC address that the switch learns and then ages when it is not in use.
- Static address—A manually entered unicast address that does not age and that is not lost when the switch resets.

The address table lists the destination MAC address, the associated VLAN ID, and port number associated with the address and the type (static or dynamic).

Address Table

With multiple MAC addresses supported on all ports, you can connect any port on the switch to individual workstations, repeaters, switches, routers, or other network devices. The switch provides dynamic addressing by learning the source address of packets it receives on each port and adding the address and its associated port number to the address table. As stations are added or removed from the network, the switch updates the address table, adding new dynamic addresses and aging out those that are not in use.

The aging interval is globally configured. However, the switch maintains an address table for each VLAN, and STP can accelerate the aging interval on a per-VLAN basis.

The switch sends packets between any combination of ports, based on the destination address of the received packet. Using the MAC address table, the switch forwards the packet only to the port associated with the destination address. If the destination address is on the port that sent the packet, the packet is filtered and not forwarded. The switch always uses the store-and-forward method: complete packets are stored and checked for errors before transmission.

MAC Addresses and VLANs

All addresses are associated with a VLAN. An address can exist in more than one VLAN and have different destinations in each. Unicast addresses, for example, could be forwarded to port 1 in VLAN 1 and ports 9, 10, and 1 in VLAN 5.

Each VLAN maintains its own logical address table. A known address in one VLAN is unknown in another until it is learned or statically associated with a port in the other VLAN.

When private VLANs are configured, address learning depends on the type of MAC address:

- Dynamic MAC addresses learned in one VLAN of a private VLAN are replicated in the associated VLANs. For example, a MAC address learned in a private-VLAN secondary VLAN is replicated in the primary VLAN.
- Static MAC addresses configured in a primary or secondary VLAN are not replicated in the associated VLANs. When you configure a static MAC address in a private VLAN primary or secondary VLAN, you should also configure the same static MAC address in all associated VLANs.

Default MAC Address Table Configuration

Table 7-2 Default MAC Address Table Configuration

Feature	Default Setting
Aging time	300 seconds
Dynamic addresses	Automatically learned
Static addresses	None configured

Address Aging Time for VLANs

Dynamic addresses are source MAC addresses that the switch learns and then ages when they are not in use. You can change the aging time setting for all VLANs or for a specified VLAN.

Setting too short an aging time can cause addresses to be prematurely removed from the table. Then when the switch receives a packet for an unknown destination, it floods the packet to all ports in the same VLAN as the receiving port. This unnecessary flooding can impact performance. Setting too long an aging time can cause the address table to be filled with unused addresses, which prevents new addresses from being learned. Flooding results, which can impact switch performance.

MAC Address Change Notification Traps

MAC address change notification tracks users on a network by storing the MAC address change activity. When the switch learns or removes a MAC address, an SNMP notification trap can be sent to the NMS. If you have many users coming and going from the network, you can set a trap-interval time to bundle the notification traps to reduce network traffic. The MAC notification history table stores MAC address activity for each port for which the trap is set. MAC address change notifications are generated for dynamic and secure MAC addresses. Notifications are not generated for self addresses, multicast addresses, or other static addresses.

Static Addresses

A static address has these characteristics:

- Is manually entered in the address table and must be manually removed.
- Can be a unicast or multicast address.
- Does not age and is retained when the switch restarts.

You can add and remove static addresses and define the forwarding behavior for them. The forwarding behavior defines how a port that receives a packet forwards it to another port for transmission. Because all ports are associated with at least one VLAN, the switch acquires the VLAN ID for the address from the ports that you specify. You can specify a different list of destination ports for each source port.

A packet with a static address that arrives on a VLAN where it has not been statically entered is flooded to all ports and not learned.

You add a static address to the address table by specifying the destination MAC unicast address and the VLAN from which it is received. Packets received with this destination address are forwarded to the interface specified with the *interface-id* option.

When you configure a static MAC address in a private-VLAN primary or secondary VLAN, you should also configure the same static MAC address in all associated VLANs. Static MAC addresses configured in a private-VLAN primary or secondary VLAN are not replicated in the associated VLAN. For more information about private VLANs, see Chapter 19, "Configuring Private VLANs."

Unicast MAC Address Filtering

When unicast MAC address filtering is enabled, the switch drops packets with specific source or destination MAC addresses. This feature is disabled by default and only supports unicast static addresses.

Follow these guidelines when using this feature:

- Multicast MAC addresses, broadcast MAC addresses, and router MAC addresses are not supported. If you specify one of these addresses when entering the **mac address-table static** *mac-addr* **vlan** *vlan-id* **drop** global configuration command, one of these messages appears:

  ```
  % Only unicast addresses can be configured to be dropped

  % CPU destined address cannot be configured as drop address
  ```

- Packets that are forwarded to the CPU are also not supported.
- If you add a unicast MAC address as a static address and configure unicast MAC address filtering, the switch either adds the MAC address as a static address or drops packets with that MAC address, depending on which command was entered last. The second command that you entered overrides the first command.

  For example, if you enter the **mac address-table static** *mac-addr* **vlan** *vlan-id* **interface** *interface-id* global configuration command followed by the **mac address-table static** *mac-addr* **vlan** *vlan-id* **drop** command, the switch drops packets with the specified MAC address as a source or destination.

  If you enter the **mac address-table static** *mac-addr* **vlan** *vlan-id* **drop** global configuration command followed by the **mac address-table static** *mac-addr* **vlan** *vlan-id* **interface** *interface-id* command, the switch adds the MAC address as a static address.

You enable unicast MAC address filtering and configure the switch to drop packets with a specific address by specifying the source or destination unicast MAC address and the VLAN from which it is received.

MAC Address Learning on a VLAN

By default, MAC address learning is enabled on all VLANs on the switch. You can control MAC address learning on a VLAN to manage the available MAC address table space by controlling which VLANs, and therefore which ports, can learn MAC addresses. Before you disable MAC address learning, be sure that you are familiar with the network topology and the switch system configuration. Disabling MAC address learning on a VLAN could cause flooding in the network.

Follow these guidelines when disabling MAC address learning on a VLAN:

- Use caution before disabling MAC address learning on a VLAN with a configured switch virtual interface (SVI). The switch then floods all IP packets in the Layer 2 domain.
- You can disable MAC address learning on a single VLAN ID (for example, **no mac address-table learning vlan 223**) or on a range of VLAN IDs (for example, **no mac address-table learning vlan 1-20, 15**).
- We recommend that you disable MAC address learning only in VLANs with two ports. If you disable MAC address learning on a VLAN with more than two ports, every packet entering the switch is flooded in that VLAN domain.
- You cannot disable MAC address learning on a VLAN that is used internally by the switch. If the VLAN ID that you enter is an internal VLAN, the switch generates an error message and rejects the command. To view internal VLANs in use, enter the **show vlan internal usage** privileged EXEC command.
- If you disable MAC address learning on a VLAN configured as a private-VLAN primary VLAN, MAC addresses are still learned on the secondary VLAN that belongs to the private VLAN and are then replicated on the primary VLAN. If you disable MAC address learning on the secondary VLAN, but not the primary VLAN of a private VLAN, MAC address learning occurs on the primary VLAN and is replicated on the secondary VLAN.
- You cannot disable MAC address learning on an RSPAN VLAN. The configuration is not allowed.
- If you disable MAC address learning on a VLAN that includes a secure port, MAC address learning is not disabled on that port. If you disable port security, the configured MAC address learning state is enabled.

To reenable MAC address learning on a VLAN, use the **default mac address-table learning vlan** *vlan-id* global configuration command. You can also reenable MAC address learning on a VLAN by entering the **mac address-table learning vlan** *vlan-id* global configuration command. The first (**default**) command returns to a default condition and therefore does not appear in the output from the **show running-config** command. The second command causes the configuration to appear in the **show running-config** privileged EXEC command display.

ARP Table Management

To communicate with a device (over Ethernet, for example), the software first must learn the 48-bit MAC address or the local data link address of that device. The process of learning the local data link address from an IP address is called *address resolution*.

The Address Resolution Protocol (ARP) associates a host IP address with the corresponding media or MAC addresses and the VLAN ID. Using an IP address, ARP finds the associated MAC address. When a MAC address is found, the IP-MAC address association is stored in an ARP cache for rapid retrieval. Then the IP datagram is encapsulated in a link-layer frame and sent over the network. Encapsulation of

IP datagrams and ARP requests and replies on IEEE 802 networks other than Ethernet is specified by the Subnetwork Access Protocol (SNAP). By default, standard Ethernet-style ARP encapsulation (represented by the **arpa** keyword) is enabled on the IP interface.

ARP entries added manually to the table do not age and must be manually removed.

How to Perform Switch Administration

Configuring Time and Date Manually

If no other source of time is available, you can manually configure the time and date after the system is restarted. The time remains accurate until the next system restart. We recommend that you use manual configuration only as a last resort. If you have an outside source to which the switch can synchronize, you do not need to manually set the system clock.

Setting the System Clock

If you have an outside source on the network that provides time services, such as an NTP server, you do not need to manually set the system clock.

Beginning in privileged EXEC mode, follow these steps to set the system clock:

	Command	Purpose
Step 1	**clock set** *hh*:*mm*:*ss day month year* or **clock set** *hh*:*mm*:*ss month day year*	Manually sets the system clock using one of these formats: • *hh*:*mm*:*ss*—Specifies the time in hours (24-hour format), minutes, and seconds. The time specified is relative to the configured time zone. • *day*—Specifies the day by date in the month. • *month*—Specifies the month by name. • *year*—Specifies the year (no abbreviation).

Configuring the Time Zone

The *minutes-offset* variable in the **clock timezone** global configuration command is available for those cases where a local time zone is a percentage of an hour different from UTC. For example, the time zone for some sections of Atlantic Canada (AST) is UTC-3.5, where the 3 means 3 hours and .5 means 50 percent. In this case, the necessary command is **clock timezone AST -3 30**.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**clock timezone** *zone hours-offset* [*minutes-offset*]	Sets the time zone. The switch keeps internal time in universal time coordinated (UTC), so this command is used only for display purposes and when the time is manually set. • *zone*—Enters the name of the time zone to be displayed when standard time is in effect. The default is UTC. • *hours-offset*—Enters the hours offset from UTC. • (Optional) *minutes-offset*—Enters the minutes offset from UTC.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring Summer Time (Daylight Saving Time)

To configure summer time (daylight saving time) in areas where it starts and ends on a particular day of the week each year, perform this task:

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**clock summer-time** *zone* **recurring** [*week day month hh*:*mm week day month hh*:*mm* [*offset*]]	Configures summer time to start and end on the specified days every year. Summer time is disabled by default. If you specify **clock summer-time** *zone* **recurring** without parameters, the summer time rules default to the United States rules. • *zone*—Specifies the name of the time zone (for example, PDT) to be displayed when summer time is in effect. • (Optional) *week*—Specifies the week of the month (1 to 5 or **last**). • (Optional) *day*—Specifies the day of the week (Sunday, Monday...). • (Optional) *month*—Specifies the month (January, February...). • (Optional) *hh*:*mm*—Specifies the time (24-hour format) in hours and minutes. • (Optional) *offset*—Specifies the number of minutes to add during summer time. The default is 60.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring Summer Time (Exact Date and Time)

To configure summer time when it does not follow a recurring pattern (configure the exact date and time of the next summer time events), perform this task:

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**clock summer-time** *zone* **date** [*month date year hh*:*mm month date year hh*:*mm* [*offset*]] or **clock summer-time** *zone* **date** [*date month year hh*:*mm date month year hh*:*mm* [*offset*]]	Configures summer time to start on the first date and end on the second date. Summer time is disabled by default. • *zone*—Specifies the name of the time zone (for example, PDT) to be displayed when summer time is in effect. • (Optional) *week*—Specifies the week of the month (1 to 5 or **last**). • (Optional) *day*—Specifies the day of the week (Sunday, Monday...). • (Optional) *month*—Specifies the month (January, February...). • (Optional) *hh*:*mm*—Specifies the time (24-hour format) in hours and minutes. • (Optional) *offset*—Specifies the number of minutes to add during summer time. The default is 60.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring a System Name

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**hostname** *name*	Manually configures a system name. The default setting is *switch*. The name must follow the rules for ARPANET hostnames. They must start with a letter, end with a letter or digit, and have as interior characters only letters, digits, and hyphens. Names can be up to 63 characters.
Step 3	**end**	Returns to privileged EXEC mode.

Setting Up DNS

If you use the switch IP address as its hostname, the IP address is used and no DNS query occurs. If you configure a hostname that contains no periods (.), a period followed by the default domain name is appended to the hostname before the DNS query is made to map the name to an IP address. The default domain name is the value set by the **ip domain-name** global configuration command. If there is a period (.) in the hostname, the Cisco IOS software looks up the IP address without appending any default domain name to the hostname.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip domain-name** *name*	Defines a default domain name that the software uses to complete unqualified hostnames (names without a dotted-decimal domain name). Do not include the initial period that separates an unqualified name from the domain name. At boot-up time, no domain name is configured; however, if the switch configuration comes from a BOOTP or Dynamic Host Configuration Protocol (DHCP) server, then the default domain name might be set by the BOOTP or DHCP server (if the servers were configured with this information).
Step 3	**ip name-server** *server-address1* [*server-address2 ... server-address6*]	Specifies the address of one or more name servers to use for name and address resolution. You can specify up to six name servers. Separate each server address with a space. The first server specified is the primary server. The switch sends DNS queries to the primary server first. If that query fails, the backup servers are queried.
Step 4	**ip domain-lookup**	(Optional) Enables DNS-based hostname-to-address translation on your switch. This feature is enabled by default. If your network devices require connectivity with devices in networks for which you do not control name assignment, you can dynamically assign device names that uniquely identify your devices by using the global Internet naming scheme (DNS).
Step 5	**end**	Returns to privileged EXEC mode.

Configuring Login Banners

Configuring a Message-of-the-Day Login Banner

You can create a single or multiline message banner that appears on the screen when someone logs in to the switch.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**banner motd** *c message c*	Specifies the message of the day. • *c*—Enters the delimiting character of your choice, for example, a pound sign (#), and press the **Return** key. The delimiting character signifies the beginning and end of the banner text. Characters after the ending delimiter are discarded. • *message*—Enters a banner message up to 255 characters. You cannot use the delimiting character in the message.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring a Login Banner

You can configure a login banner to be displayed on all connected terminals. This banner appears after the MOTD banner and before the login prompt.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**banner login** *c message c*	Specifies the login message. • *c*—Enters the delimiting character of your choice, for example, a pound sign (#), and press the **Return** key. The delimiting character signifies the beginning and end of the banner text. Characters after the ending delimiter are discarded. • *message*—Enters a login message up to 255 characters. You cannot use the delimiting character in the message.
Step 3	**end**	Returns to privileged EXEC mode.

Managing the MAC Address Table

Changing the Address Aging Time

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mac address-table aging-time** [**0** | *10-1000000*] [**vlan** *vlan-id*]	Sets the length of time that a dynamic entry remains in the MAC address table after the entry is used or updated. The range is 10 to 1000000 seconds. The default is 300. You can also enter 0, which disables aging. Static address entries are never aged or removed from the table. • *vlan-id*—Valid IDs are 1 to 4096.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring MAC Address Change Notification Traps

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**snmp-server host** *host-addr* {**traps** | **informs**} {**version** {**1** | **2c** | **3**}} *community-string notification-type*	Specifies the recipient of the trap message. • *host-addr*—Specifies the name or address of the NMS. • **traps** (the default)—Sends SNMP traps to the host. • **informs**—Sends SNMP informs to the host. • Specifies the SNMP version to support. Version 1, the default, is not available with informs. • *community-string*—Specifies the string to send with the notification operation. You can set this string by using the **snmp-server host** command, but we recommend that you define this string by using the **snmp-server community** command before using the **snmp-server host** command. • *notification-type*—Uses the **mac-notification** keyword.
Step 3	**snmp-server enable traps mac-notification change**	Enables the switch to send MAC address change notification traps to the NMS.
Step 4	**mac address-table notification change**	Enables the MAC address change notification feature.
Step 5	**mac address-table notification change** [**interval** *value*] [**history-size** *value*]	Enters the trap interval time and the history table size. • (Optional) **interval** *value*—Specifies the notification trap interval in seconds between each set of traps that are generated to the NMS. The range is 0 to 2147483647 seconds; the default is 1 second. • (Optional) **history-size** *value*—Specifies the maximum number of entries in the MAC notification history table. The range is 0 to 500; the default is 1.
Step 6	**interface** *interface-id*	Enters interface configuration mode, and specifies the Layer 2 interface on which to enable the SNMP MAC address notification trap.

	Command	Purpose
Step 7	**snmp trap mac-notification change** {**added** | **removed**}	Enables the MAC address change notification trap on the interface. • Enables the trap when a MAC address is **added** on this interface. • Enables the trap when a MAC address is **removed** from this interface.
Step 8	**end**	Returns to privileged EXEC mode.

Configuring MAC Address Move Notification Traps

When you configure MAC-move notification, an SNMP notification is generated and sent to the network management system whenever a MAC address moves from one port to another within the same VLAN.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**snmp-server host** *host-addr* {**traps** | **informs**} {**version** {**1** | **2c** | **3**}} *community-string notification-type*	Specifies the recipient of the trap message. • *host-addr*—Specifies the name or address of the NMS. • **traps** (the default)—Sends SNMP traps to the host. • **informs**—Sends SNMP informs to the host. • **version**—Specifies the SNMP version to support. Version 1, the default, is not available with informs. • *community-string*—Specifies the string to send with the notification operation. You can set this string by using the **snmp-server host** command, but we recommend that you define this string by using the **snmp-server community** command before using the **snmp-server host** command. • *notification-type*—Uses the **mac-notification** keyword.
Step 3	**snmp-server enable traps mac-notification move**	Enables the switch to send MAC address move notification traps to the NMS.
Step 4	**mac address-table notification mac-move**	Enables the MAC address move notification feature.
Step 5	**end**	Returns to privileged EXEC mode.

Configuring MAC Threshold Notification Traps

When you configure MAC threshold notification, an SNMP notification is generated and sent to the network management system when a MAC address table threshold limit is reached or exceeded.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**snmp-server host** *host-addr* {**traps** | **informs**} {**version** {**1** | **2c** | **3**}} *community-string notification-type*	Specifies the recipient of the trap message. • *host-addr*—Specifies the name or address of the NMS. • **traps** (the default)—Sends SNMP traps to the host. • **informs**—Sends SNMP informs to the host. • **version**—Specifies the SNMP version to support. Version 1, the default, is not available with informs. • *community-string*—Specifies the string to send with the notification operation. You can set this string by using the **snmp-server host** command, but we recommend that you define this string by using the **snmp-server community** command before using the **snmp-server host** command. • *notification-type*—Uses the **mac-notification** keyword.
Step 3	**snmp-server enable traps mac-notification threshold**	Enables the switch to send MAC threshold notification traps to the NMS.
Step 4	**mac address-table notification threshold**	Enables the MAC address threshold notification feature.
Step 5	**mac address-table notification threshold** [**limit** *percentage*] | [**interval** *time*]	Enters the threshold value for the MAC address threshold usage monitoring. • (Optional) **limit** *percentage*—Specifies the percentage of the MAC address table use; valid values are from 1 to 100 percent. The default is 50 percent. • (Optional) **interval** *time*—Specifies the time between notifications; valid values are greater than or equal to 120 seconds. The default is 120 seconds.
Step 6	**end**	Returns to privileged EXEC mode.

Adding and Removing Static Address Entries

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mac address-table static** *mac-addr* **vlan** *vlan-id* **interface** *interface-id*	Adds a static address to the MAC address table. • *mac-addr*—Specifies the destination MAC unicast address to add to the address table. Packets with this destination address received in the specified VLAN are forwarded to the specified interface. • *vlan-id*—Specifies the VLAN for which the packet with the specified MAC address is received. Valid VLAN IDs are 1 to 4096. • *interface-id*—Specifies the interface to which the received packet is forwarded. Valid interfaces include physical ports or port channels. For static multicast addresses, you can enter multiple interface IDs. For static unicast addresses, you can enter only one interface at a time, but you can enter the command multiple times with the same MAC address and VLAN ID.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring Unicast MAC Address Filtering

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mac address-table static** *mac-addr* **vlan** *vlan-id* **drop**	Enables unicast MAC address filtering and configures the switch to drop a packet with the specified source or destination unicast static address. • *mac-addr*—Specifies a source or destination unicast MAC address. Packets with this MAC address are dropped. • *vlan-id*—Specifies the VLAN for which the packet with the specified MAC address is received. Valid VLAN IDs are 1 to 4096.
Step 3	**end**	Returns to privileged EXEC mode.

Disabling MAC Address Learning on a VLAN

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**no mac address-table learning vlan** *vlan-id*	Disables MAC address learning on the specified VLAN or VLANs. You can specify a single VLAN ID or a range of VLAN IDs separated by a hyphen or comma. Valid VLAN IDs are 1 to 4096.
Step 3	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining Switch Administration

Command	Purpose
clear mac address-table dynamic	Removes all dynamic entries.
clear mac address-table dynamic address *mac-address*	Removes a specific MAC address.
clear mac address-table dynamic interface *interface-id*	Removes all addresses on the specified physical port or port channel.
clear mac address-table dynamic vlan *vlan-id*	Removes all addresses on a specified VLAN.
show clock [**detail**]	Displays the time and date configuration.
show ip igmp snooping groups	Displays the Layer 2 multicast entries for all VLANs or the specified VLAN.
show mac address-table address	Displays MAC address table information for the specified MAC address.
show mac address-table aging-time	Displays the aging time in all VLANs or the specified VLAN.
show mac address-table count	Displays the number of addresses present in all VLANs or the specified VLAN.
show mac address-table dynamic	Displays only dynamic MAC address table entries.
show mac address-table interface	Displays the MAC address table information for the specified interface.
show mac address-table learning	Displays MAC address learning status of all VLANs or the specified VLAN.
show mac address-table notification	Displays the MAC notification parameters and history table.
show mac address-table static	Displays only static MAC address table entries.
show mac address-table vlan	Displays the MAC address table information for the specified VLAN.

Configuration Examples for Performing Switch Admininistration

Setting the System Clock: Example

This example shows how to manually set the system clock to 1:32 p.m. on July 23, 2001:

```
Switch# clock set 13:32:00 23 July 2001
```

Configuring Summer Time: Examples

The first part of the **clock summer-time** global configuration command specifies when summer time begins, and the second part specifies when it ends. All times are relative to the local time zone. The start time is relative to standard time. The end time is relative to summer time. If the starting month is after the ending month, the system assumes that you are in the southern hemisphere.

This example (for daylight savings time) shows how to specify that summer time starts on the first Sunday in April at 02:00 and ends on the last Sunday in October at 02:00:

```
Switch(config)# clock summer-time PDT recurring 1 Sunday April 2:00 last Sunday October
2:00
```

This example shows how to set summer time to start on October 12, 2000, at 02:00, and end on April 26, 2001, at 02:00:

```
Switch(config)# clock summer-time pdt date 12 October 2000 2:00 26 April 2001 2:00
```

Configuring a MOTD Banner: Examples

This example shows how to configure a MOTD banner for the switch by using the pound sign (#) symbol as the beginning and ending delimiter:

```
Switch(config)# banner motd #
This is a secure site. Only authorized users are allowed.
For access, contact technical support.
#
Switch(config)#
```

This example shows the banner that appears from the previous configuration:

```
Unix> telnet 172.2.5.4
Trying 172.2.5.4...
Connected to 172.2.5.4.
Escape character is '^]'.

This is a secure site. Only authorized users are allowed.
For access, contact technical support.

User Access Verification

Password:
```

Configuring a Login Banner: Example

This example shows how to configure a login banner for the switch by using the dollar sign ($) symbol as the beginning and ending delimiter:

```
Switch(config)# banner login $
Access for authorized users only. Please enter your username and password.
$
Switch(config)#
```

Configuring MAC Address Change Notification Traps: Example

This example shows how to specify 172.20.10.10 as the NMS, enable the switch to send MAC address notification traps to the NMS, enable the MAC address-change notification feature, set the interval time to 123 seconds, set the history-size to 100 entries, and enable traps whenever a MAC address is added on the specified port.

```
Switch(config)# snmp-server host 172.20.10.10 traps private mac-notification
Switch(config)# snmp-server enable traps mac-notification change
Switch(config)# mac address-table notification change
Switch(config)# mac address-table notification change interval 123
Switch(config)# mac address-table notification change history-size 100
```

```
Switch(config)# interface gigabitethernet1/2

Switch(config-if)# snmp trap mac-notification change added
```

Sending MAC Address Move Notification Traps: Example

This example shows how to specify 172.20.10.10 as the NMS, enable the switch to send MAC address move notification traps to the NMS, enable the MAC address move notification feature, and enable traps when a MAC address moves from one port to another.

```
Switch(config)# snmp-server host 172.20.10.10 traps private mac-notification
Switch(config)# snmp-server enable traps mac-notification move
Switch(config)# mac address-table notification mac-move
```

Configuring MAC Threshold Notification Traps: Example

This example shows how to specify 172.20.10.10 as the NMS, enable the MAC address threshold notification feature, set the interval time to 123 seconds, and set the limit to 78 per cent.

```
Switch(config)# snmp-server host 172.20.10.10 traps private mac-notification
Switch(config)# snmp-server enable traps mac-notification threshold
Switch(config)# mac address-table notification threshold
Switch(config)# mac address-table notification threshold interval 123
Switch(config)# mac address-table notification threshold limit 78
```

Adding the Static Address to the MAC Address Table: Example

This example shows how to add the static address c2f3.220a.12f4 to the MAC address table. When a packet is received in VLAN 4 with this MAC address as its destination address, the packet is forwarded to the specified port:

```
Switch(config)# mac address-table static c2f3.220a.12f4 vlan 4 interface
gigabitethernet1/1
```

Configuring Unicast MAC Address Filtering: Example

This example shows how to enable unicast MAC address filtering and to configure the switch to drop packets that have a source or destination address of c2f3.220a.12f4. When a packet is received in VLAN 4 with this MAC address as its source or destination, the packet is dropped:

```
Switch(config)# mac address-table static c2f3.220a.12f4 vlan 4 drop
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
Cisco IOS routing commands.	*Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 8

Configuring PTP

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for Configuring PTP

- To use this feature, the switch must be PTP-capable. Refer to your switch release notes.

Restrictions for Configuring PTP

- To use this feature, the switch must be running the LAN Base image.

Information About Configuring PTP

Precision Time Protocol

The IEEE 1588 standard describes the use of PTP for fault-tolerant synchronization of network real-time clocks.

The clocks in a PTP network are organized into a master-slave hierarchy. The grandmaster clock is called the Best Master Clock (BMC), and is the root of the master-slave clock hierarchy. PTP uses the BMC algorithm to identify the master clock for synchronization.

The master clock is a time source on the network that can be synchronized to a highly accurate time source such as a Global Positioning System (GPS) clock. The slaves are the other network devices that synchronize their clocks to the master clock. The parent is the clock to which the member-slave clocks synchronize. Timing messages between the master and slave clocks ensure continued synchronization.

Synchronization behavior depends on the PTP clock setting mode that you configure on the switch. The mode can be boundary, end-to-end transparent, or forward:

- A switch clock in boundary mode participates in the selection of the most accurate master clock. If more accurate clocks are not detected, that switch clock becomes the master clock. If a more accurate clock is found among the slave clocks, then the switch synchronizes to that clock and becomes a slave clock. After initial synchronization, the switch and the connected devices exchange timing messages to correct the changes caused by clock offsets and network delays.
- A switch clock in end-to-end transparent mode synchronizes all switch ports with the master clock. This switch does not participate in master clock selection and uses the default PTP clock mode on all ports.
- A switch clock in forward mode allows incoming PTP packets to pass-through the switch as normal multicast traffic.

When the switch is in PTP forward mode, PTP configuration is not available except when changing PTP mode to another mode. You can only configure per-port PTP when the switch is in boundary mode.

How to Configure PTP

- Default PTP Settings, page 8-2
- Setting Up PTP, page 8-3

Default PTP Settings

By default, PTP is enabled on all the Fast Ethernet and Gigabit Ethernet ports on the base switch module. The default PTP mode on all ports is end-to-end transparent.

Table 8-1 Default PTP Settings

Feature	Default Setting
PTP boundary mode	Disabled.
PTP forward mode	Disabled.
PTP end-to-end transparent mode	Enabled.
PTP priority 1 and PTP priority 2	Default priority number is 128.
PTP announce interval	2 seconds.
PTP announce receipt time out	3 messages.
PTP delay request interval	32 seconds.
PTP sync interval	1 second.
PTP sync limit	500000000 nanoseconds.

Setting Up PTP

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Enters interface configuration mode.
Step 3	**ptp** {**announce** {**interval** *value* | **timeout** *value*} | **delay-req interval** *value* | **enable** | **sync** {**interval** *value* | **limit** *value*}}	Specifies the settings for the timing messages. These options are available only when the switch is in boundary mode. • **announce interval** *value*—Sets the time to send announce messages. The range is 0 to 4 seconds. The default is 1 (2 seconds). • **announce timeout** *value*— Sets the time to announce timeout messages. The range is 2 to 10 seconds. The default is 3 (8 seconds). • **delay-req interval** *value*—Sets the time for slave devices to send delay request messages when the port is in the master clock state. The range is -1 second to 6 seconds. The default is 5 (32 seconds). • **enable**—Enables PTP on the port base module. • **sync interval** *value*—Sets the time to send synchronization messages. The range is –1 second to 1 second. The default is 1 second. • **sync limit** *value*—Sets the maximum clock offset value before PTP attempts to resynchronize. The range is from 50 to 500000000 nanoseconds. The default is 500000000 nanoseconds.
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show running-config**	Verifies your entries.
Step 6	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Monitoring and Maintaining the PTP Configuration

Table 8-2 Commands for Displaying the PTP Configuration

Command	Purpose
show ptp clock	Displays the PTP clock properties.
show ptp foreign-master-record	Displays the PTP foreign master data set.
show ptp parent	Displays the parent and grandmaster clock properties.
show ptp port	Displays all the PTP port properties.
show ptp port FastEthernet *interface*	Displays the PTP FastEthernet properties on the specified port.
show ptp port GigabitEthernet *interface*	Displays the PTP Gigabit Ethernet properties on the specified port.
show ptp time-property	Displays the PTP time properties.

Troubleshooting the PTP Configuration

Table 8-3 Commands for Troubleshooting the PTP Configuration

Command	Purpose
debug ptp bmc	Enables debugging of the PTP Best Master Clock Algorithm.
debug ptp clock-correction	Enables debugging of PTP clock correction.
debug ptp collision	Enables debugging of PTP source collision.
debug ptp error	Enables debugging of PTP errors.
debug ptp event	Enables debugging of PTP state event.
debug ptp messages	Enables debugging of PTP messages.
debug ptp transparent-clock	Enables debugging of the PTP transparent clock.

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 9

Configuring PROFINET

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Restrictions for Configuring PROFINET

The switch does not support isochronous real-time communication channels.

Information About Configuring PROFINET

PROFINET is the PROFIBUS International (PI) open Industrial Ethernet Standard that uses TCP/IP and IT standards for automation control. PROFINET is particularly useful for industrial automation systems and process control networks, in which motion control and precision control of instrumentation and test equipment are important. It emphasizes data exchange and defines communication paths to meet speed requirements. PROFINET communication is scalable on three levels:

- Normal non-real-time communication uses TCP/IP and enables bus cycle times of approximately 100 ms.
- Real-time communication enables cycle times of approximately 10 ms.
- Isochronous real-time communication enables cycle times of approximately 1 ms.

PROFINET I/O is a modular communication framework for distributed automation applications. PROFINET I/O uses cyclic data transfer to exchange data, alarms, and diagnostic information with programmable controllers, input/output (I/O) devices, and other automation controllers (for example, motion controllers).

PROFINET I/O recognizes three classes of devices:

- I/O devices
- I/O controllers
- I/O supervisors

PROFINET Device Roles

Figure 9-1 PROFINET Device Roles

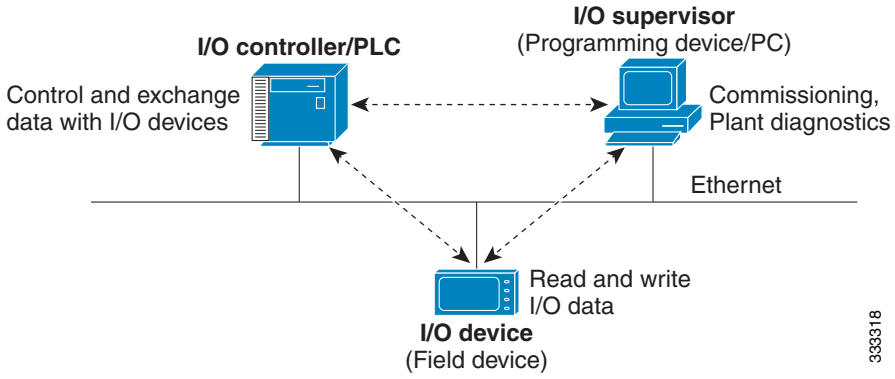

An I/O controller is a programmable logic controller (PLC) that controls I/O devices and exchanges data such as configuration, alarms, and I/O data through an automation program. The I/O controller and the I/O supervisor exchange diagnostic information. The I/O controller shares configuration and input/output information with the I/O device and receives alarms from the I/O device.

PROFINET is designed to be the sole or primary management system platform. Because the I/O controller detects the switch with the Discovery and Configuration Protocol (DCP), and sets the device name and IP address, you do not need to enter Cisco IOS commands for the basic configuration. For advanced configurations (for example, QoS, DHCP, and similar features) you must use Cisco IOS commands on the switch because these features cannot be configured by using PROFINET.

An I/O supervisor is an engineering station, such as a human machine interface (HMI) or PC, used for commissioning, monitoring, and diagnostic analysis. The I/O supervisor exchanges diagnostic, status, control, and parameter information with the I/O device.

An I/O device is a distributed input/output device such as a sensor, an actuator, or a motion controller.

Note The switch acts as an I/O device, providing a PROFINET management connection to the I/O controllers.

In a PROFINET I/O system, all the I/O devices communicate over an Ethernet communication network to meet the automation industry requirement for bus cycle times of less than 100 ms. The network uses switches and full-duplex data exchange to avoid data collisions.

PROFINET Device Data Exchange

After PROFINET uses DCP to discover devices, including the switch, they establish application relationships (ARs) and communication relationships (CRs). After a connection is established and information about device parameters is exchanged, input and output data is exchanged. The switch uses non-real-time CRs to exchange the data attributes listed in Table 9-1 and Table 9-2.

Table 9-1 PROFINET I/O Switch Attributes

PROFINET I/O Switch Configuration Attributes	Value or Action
Device name	Configures a name for the device.
TCP/IP	IP address, subnet mask, default gateway, SVI.
Primary temperature alarm	Enables or disables monitoring for the specified alarm.
Secondary temperature alarm	Enables or disables monitoring for the specified alarm.
RPS failed alarm	Enables or disables monitoring for the specified alarm.
Relay major alarm	Enables or disables monitoring for the specified alarm.
Reset to factory defaults	Uses the PROFINET I/O controller to reset the switch to factory defaults. This action removes the startup configuration and reloads the switch.
Relay major configuration	Specifies the type of port alarm (for example, link fault) that triggers the major relay. Any port configured with the specified alarm type can trigger the major relay.

Table 9-2 PROFINET I/O Port Attributes

PROFINET I/O Port Configuration Attributes	Value or Action
Speed	10/100/1000/auto,
Duplex	Half/full/auto,
Port mode	Access/trunk,
Link status	Shut down/no shut down,
Configure rate limiting	Broadcast, unicast, multicast threshold exceeds configured levels.
Port link fault alarm	Enables or disables monitoring for specified alarm.
Port not forwarding alarm	Enables or disables monitoring for specified alarm.
Port not operating alarm	Enables or disables monitoring for specified alarm.
Port FCS threshold alarm	Enables or disables monitoring for specified alarm.

PROFINET devices are integrated by using a general station description (GSD) file that contains the data for engineering and data exchange between the I/O controller, the I/O supervisor, and the I/O devices, including the switch. Each PROFINET I/O field device must have an associated GSD file that describes the properties of the device and contains all this information required for configuration:

- Device identification information (device ID, vendor ID and name, product family, number of ports)
- Number and types of pluggable modules
- The Cisco IE 2000 8-port expander modules are not hot-swappable. Turn off the switch before connecting or disconnecting expander modules.
- Error text for diagnostic information
- Communication parameters for I/O devices, including the minimum cycle time, the reduction ratio, and the watch dog time

Note Although the Cisco IE 2000 switch has a default reduction ratio of 128 ms, we recommend a reduction ratio of 256 ms or 512 ms to reduce the load on the switch CPU when the switch uses a complex configuration.

- Configuration data for the I/O device modules, including speed, duplex, VLAN, port security information, alarms, and broadcast-rate-limiting thresholds
- Parameters configured for I/O device modules for the attributes listed in Table 9-2

The GSD file is on the switch, but the I/O supervisor uses this file.

Note You must use the GSD file that is associated with the Cisco IOS release on the switch to manage your PROFINET network. Both the I/O supervisor and the Cisco IOS software alert you to a mismatch between the GSD file and the switch Cisco IOS software version.

How to Configure PROFINET

Configuring PROFINET

You can use either the PROFINET software on the I/O supervisor or the Cisco IOS software for basic switch configuration.

Default Configuration

PROFINET is enabled by default on all the base switch module and expansion-unit Ethernet ports. If PROFINET has been disabled, follow the instructions in the "Enabling PROFINET" section on page 9-4.

Enabling PROFINET

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**profinet**	Enables PROFINET on the switch.
Step 3	**profinet id** *line*	(Optional) Sets the PROFINET device identifier (ID) by using the Cisco IOS software. The maximum length is 240 characters. The only special characters allowed are the period (.) and hyphen (-), and they are allowed only in specific positions within the ID string. It can have multiple labels within the string. Each label can be from 1 to 63 characters, and labels must be separated by a period (.). The final character in the string must not be zero (0). For more details about configuring the PROFINET ID, see the PROFINET specification, document number TC2-06-0007a, filename PN-AL-protocol_2722_V22_Oct07, available from PROFIBUS.

	Command	Purpose
Step 4	**profinet vlan** *vlan id*	(Optional) Changes the VLAN number. The default VLAN number is 1. The VLAN ID range is 1-4096.
Step 5	**end**	Returns to privileged EXEC mode.
Step 6	**show running-config**	Verifies your entries.
Step 7	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Monitoring and Maintaining PROFINET

Table 9-3 Commands for Displaying the PROFINET Configuration

Command	Purpose
show profinet sessions	Displays the currently connected PROFINET sessions.
show profinet status	Displays the status of the PROFINET subsystem.

Troubleshooting PROFINET

The PLC has LEDs that display red for alarms, and the I/O supervisor software monitors those alarms.

To troubleshoot PROFINET use the **debug profinet** privileged EXEC command with the keywords shown in Table 9-4. Be aware that the output of a **debug** command might cause a serial link to fail. You should use these commands only under the guidance of a Cisco Technical Support engineer. When you use this command, use Telnet to access the Cisco IOS command-line interface (CLI) by using Ethernet rather than a serial port.

Table 9-4 Commands for Troubleshooting the PROFINET Configuration

Command	Purpose
debug profinet alarm	Displays the alarm status (on or off) and content of PROFINET alarms.
debug profinet cyclic	Displays information about the time-cycle-based PROFINET Ethernet frames.
debug profinet error	Displays the PROFINET session errors.
debug profinet packet ethernet	Displays information about the PROFINET Ethernet packets.
debug profinet packet udp	Displays information about the PROFINET Upper Layer Data Protocol (UDP) packets.
debug profinet platform	Displays information about the interaction between the Cisco IOS software and PROFINET.
debug profinet topology	Displays the PROFINET topology packets received.
debug profinet trace	Displays a group of traced debug output logs.

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 10

Configuring CIP

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Restrictions for Configuring CIP

CIP can be enabled on only one VLAN on the switch.

Information About Configuring CIP

The Common Industrial Protocol (CIP) is an industrial protocol for industrial automation applications. It is supported by Open DeviceNet Vendors Association (ODVA), an organization that supports network technologies based upon CIP such as DeviceNet, EtherNet/IP, CIP Safety and CIP Sync.

Previously known as Control and Information Protocol, CIP encompasses a comprehensive suite of messages and services for the collection of manufacturing automation applications - control, safety, synchronization, motion, configuration and information. CIP allows users to integrate these manufacturing applications with enterprise-level Ethernet networks and the Internet.

How to Configure CIP

Default Configuration

By default, CIP is not enabled.

Enabling CIP

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**cip security** {**password** *password* | **window timeout** *value*}	Sets CIP security options on the switch.
Step 3	**interface vlan 20**	Enters interface configuration mode.
Step 4	**cip enable**	Enables CIP on a VLAN.
Step 5	**end**	Returns to privileged EXEC mode.
Step 6	**show running-config**	Verifies your entries.
Step 7	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Monitoring CIP

Table 10-1 Commands for Displaying the CIP Configuration

Command	Purpose
show cip {**connection** | **faults** | **file** | **miscellaneous** | **object** | **security**| **session** | **status**}	Displays information about the CIP subsystem.

Troubleshooting CIP

Table 10-2 Commands for Troubleshooting the CIP Configuration

Command	Purpose
debug cip {**assembly** | **connection manager** | **errors** | **event** | **file** | **io** | **packet** | **request response** | **security** | **session** | **socket**}	Enables debugging of the CIP subsystem.

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
CIP configuration through Express Setup	*Cisco IE 2000 Switch Getting Started Guide*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 11

Configuring SDM Templates

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for Configuring SDM Templates

You must enter the **reload** privileged EXEC command to have your configured SDM template take effect.

Restrictions for Configuring SDM Templates

- For IPv6 routing support, you must be running the LAN Base image on the switch.
- When you select and configure SDM templates, you must reload the switch for the configuration to take effect.
- If you try to configure IPv6 features without first selecting a dual IPv4 and IPv6 template, a warning message is generated.
- Using the dual-stack templates results in less TCAM capacity allowed for each resource, so do not use if you plan to forward only IPv4 traffic.

Information About Configuring SDM Templates

SDM Templates

You can use SDM templates to configure system resources in the switch to optimize support for specific features, depending on how the switch is used in the network.

You can select a template to provide maximum system usage for some functions or use the default template to balance resources.

To allocate ternary content addressable memory (TCAM) resources for different usages, the switch SDM templates prioritize system resources to optimize support for certain features. When running the LAN Base image, you can select SDM templates to optimize these features:

- Default—The default template gives balance to all Layer 2 functions.
- Dual IPv6 and IPv6—Allows the switch to be used in dual-stack environments (supporting both IPv4 and IPv6).
- LAN Base Routing—The routing template maximizes system resources for IPv4 unicast routing, typically required for a router or aggregator in the center of a network.

See the "Dual IPv4 and IPv6 SDM Default Template" section on page 11-3.

Note A switch running the LAN Lite image supports only the default SDM template.

Table 11-1 Approximate Number of Feature Resources Allowed by IPv4 Templates

Resource	Default
Unicast MAC addresses	12 K
Internet Group Management Protocol (IGMP) groups and multicast routes	1 K
IPv4 unicast routes	0
Policy-based routing access control entries (ACEs)	0
IPv4 or MAC QoS ACEs	0.75 K
IPv4 or MAC security ACEs	1 K

Table 11-2 Approximate Number of Feature Resources Allowed by Each Template

Resource	Default	QoS	Routing
Unicast MAC addresses	8 K	8 K	2 K
IGMP groups and multicast routes	256	256	1 K
Unicast routes	0		4 K
• Directly connected hosts	0		2 K
• Indirect routes	0		2 K
Policy-based routing ACEs	0		512
QoS classification ACEs	375	625	625
Security ACEs	375	125	375 K
Layer 2 VLANs	1 K	1 K	1 K

The first eight rows in the tables (unicast MAC addresses through security ACEs) represent approximate hardware boundaries set when a template is selected. If a section of a hardware resource is full, all processing overflow is sent to the CPU, seriously impacting switch performance. The last row is a guideline used to calculate hardware resource consumption related to the number of Layer 2 VLANs on the switch.

Dual IPv4 and IPv6 SDM Default Template

You can select an SDM template to support IP Version 6 (IPv6) switching. For more information about IPv6 and how to configure IPv6 routing, see Chapter 41, "Configuring Static IP Unicast Routing."

This software release does not support Policy-Based Routing (PBR) when forwarding IPv6 traffic. The software supports IPv4 PBR only when the **dual-ipv4-and-ipv6 routing** template is configured.

The dual IPv4 and IPv6 template allows the switch to be used in dual-stack environments (supporting both IPv4 and IPv6). Using the dual-stack templates results in less TCAM capacity allowed for each resource. You should not use this template if you plan to forward only IPv4 traffic.

These SDM templates support IPv4 and IPv6 environments:

- Dual IPv4 and IPv6 default template—Supports Layer 2, QoS, and ACLs for IPv4; and Layer 2, IPv6 host, and ACLs for IPv6.
- Dual IPv4 and IPv6 routing template—Supports Layer 2, multicast, routing (including policy-based routing), QoS, and ACLs for IPv4; and Layer 2, routing, and ACLs for IPv6.

Note An IPv4 route requires only one TCAM entry. Because of the hardware compression scheme used for IPv6, an IPv6 route can take more than one TCAM entry, reducing the number of entries forwarded in hardware. For example, for IPv6 directly connected IP addresses, the desktop template might allow less than two thousand entries.

Table 11-3 Approximate Feature Resources Allowed by Dual IPv6-IPv6 Templates[1]

Resource	IPv4-and-IPv6 Default	IPv4-and-IPv6 Routing
Unicast MAC addresses	8 K	1K
IPv4 IGMP groups and multicast routes	0.25 K	0. 5 K
Total IPv4 unicast routes:	0	2 K
• Directly connected IPv4 hosts	0	1 K
• Indirect IPv4 routes	0	1 K
IPv6 multicast groups	0.375 K	0.625 K
Total IPv6 unicast routes:	0	1.375 K
• Directly connected IPv6 addresses	0	1 K
• Indirect IPv6 unicast routes	0	0.375 K
IPv4 policy-based routing ACEs	0	0.125 K
IPv4 or MAC QoS ACEs (total)	0.375 K	0.375 K
IPv4 or MAC security ACEs (total)	0.375 K	0.125 K
IPv6 policy-based routing ACEs[2]	0	0.125 K

Table 11-3 Approximate Feature Resources Allowed by Dual IPv6-IPv6 Templates[1] (continued)

Resource	IPv4-and-IPv6 Default	IPv4-and-IPv6 Routing
IPv6 QoS ACEs	0	0.125 K
IPv6 security ACEs	0.125 K	0.125 K

1. Template estimates are based on a switch with 8 routed interfaces and approximately 1000 VLANs.
2. IPv6 policy-based routing is not supported.

How to Configure the Switch SDM Templates

Setting the SDM Template

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**sdm prefer** {**default** | **dual-ipv4-and-ipv6** {**default**} | **lanbase-routing**}	Specifies the SDM template to be used on the switch: • **default**—Gives balance to all functions. • **dual-ipv4-and-ipv6**—Selects a template that supports both IPv4 and IPv6 routing. – **default**—Balances IPv4 and IPv6 Layer 2 functionality. • **lanbase-routing**—Maximizes IPv4 routing on the switch. Use the **no sdm prefer** command to set the switch to the default template. The default template balances the use of system resources.
Step 3	**end**	Returns to privileged EXEC mode.
Step 4	**reload**	Reloads the operating system.

Monitoring and Maintaining SDM Templates

This is an example of output from the **show sdm prefer default** command:

```
Switch# show sdm prefer default
 "default" template:
 The selected template optimizes the resources in
 the switch to support this level of features for
 0 routed interfaces and 1024 VLANs.

  number of unicast mac addresses:                  8K
  number of IPv4 IGMP groups:                       0.25K
  number of IPv4/MAC qos aces:                      0.375k
  number of IPv4/MAC security aces:                 0.375k
```

This is an example of output from the **show sdm prefer dual-ipv4-and-ipv6 default** command:

```
Switch# show sdm prefer dual-ipv4-and-ipv6 default
 "dual-ipv4-and-ipv6 default" template:
 The selected template optimizes the resources in
 the switch to support this level of features for
 0 routed interfaces and 1024 VLANs.

  number of unicast mac addresses:                  7.5K
  number of IPv4 IGMP groups + multicast routes:    0.25K
  number of IPv4 unicast routes:                    0
  number of IPv6 multicast groups:                  0.375k
  number of directly-connected IPv6 addresses:      0
  number of indirect IPv6 unicast routes:           0
  number of IPv4 policy based routing aces:         0
  number of IPv4/MAC qos aces:                      0.375k
  number of IPv4/MAC security aces:                 0.375k
  number of IPv6 policy based routing aces:         0
  number of IPv6 qos aces:                          0
  number of IPv6 security aces:                     0.125k
```

This is an example of output from the **show sdm prefer lanbase-routing** command:

```
Switch# show sdm prefer lanbase-routing
 "lanbase-routing" template:
 The selected template optimizes the resources in
 the switch to support this level of features for
 8 routed interfaces and 1005 VLANs.

  number of unicast mac addresses:                  4K
  number of IPv4 IGMP groups + multicast routes:    0.25K
  number of IPv4 unicast routes:                    4.25K
    number of directly-connected IPv4 hosts:        4K
    number of indirect IPv4 routes:                 0.25K
  number of IPv4 policy based routing aces:         0
  number of IPv4/MAC qos aces:                      0.375k
  number of IPv4/MAC security aces:                 0.375k
```

Configuration Examples for Configuring SDM Templates

Configuring the IPv4-and-IPv6 Default Template: Example

This example shows how to configure the IPv4-and-IPv6 default template on a desktop switch:

```
Switch(config)# sdm prefer dual-ipv4-and-ipv6 default
Switch(config)# exit
Switch# reload
Proceed with reload? [confirm]
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 12

Configuring Switch-Based Authentication

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for Configuring Switch-Based Authentication

- If you configure an SDM template and then perform the **show sdm prefer** command, the template currently in use displays.
- You must enter the **reload** privileged EXEC command to have your configured SDM template take effect.
- You should have access to and should configure a RADIUS server before configuring RADIUS features on your switch.
- At a minimum, you must identify the host or hosts that run the RADIUS server software and define the method lists for RADIUS authentication. You can optionally define method lists for RADIUS authorization and accounting.

Restrictions for Configuring Switch-Based Authentication

- To use the Radius CoA interface, a session must already exist on the switch. CoA can be used to identify a session and enforce a disconnect request. The update affects only the specified session.
- To use Secure Shell, you must install the cryptographic (encrypted) software image on your switch. You must obtain authorization to use this feature and to download the cryptographic software files from Cisco.com. For more information, see the release notes for this release.

Information About Configuring Switch-Based Authentication

Prevention for Unauthorized Switch Access

You can prevent unauthorized users from reconfiguring your switch and viewing configuration information. Typically, you want network administrators to have access to your switch while you restrict access to users who dial from outside the network through an asynchronous port, connect from outside the network through a serial port, or connect through a terminal or workstation from within the local network.

To prevent unauthorized access into your switch, you should configure one or more of these security features:

- At a minimum, you should configure passwords and privileges at each switch port. These passwords are locally stored on the switch. When users attempt to access the switch through a port or line, they must enter the password specified for the port or line before they can access the switch.
- For an additional layer of security, you can also configure username and password pairs, which are locally stored on the switch. These pairs are assigned to lines or ports and authenticate each user before that user can access the switch. If you have defined privilege levels, you can also assign a specific privilege level (with associated rights and privileges) to each username and password pair.
- If you want to use username and password pairs, but you want to store them centrally on a server instead of locally, you can store them in a database on a security server. Multiple networking devices can then use the same database to obtain user authentication (and, if necessary, authorization) information.
- You can also enable the login enhancements feature, which logs both failed and unsuccessful login attempts. Login enhancements can also be configured to block future login attempts after a set number of unsuccessful attempts are made.

Password Protection

A simple way of providing terminal access control in your network is to use passwords and assign privilege levels. Password protection restricts access to a network or network device. Privilege levels define what commands users can enter after they have logged into a network device.

Default Password and Privilege Level Configuration

Table 12-1 Default Password and Privilege Levels

Feature	Default Setting
Enable password and privilege level	No password is defined. The default is level 15 (privileged EXEC level). The password is not encrypted in the configuration file.
Enable secret password and privilege level	No password is defined. The default is level 15 (privileged EXEC level). The password is encrypted before it is written to the configuration file.
Line password	No password is defined.

Enable Secret Passwords with Encryption

To provide an additional layer of security, particularly for passwords that cross the network or that are stored on a Trivial File Transfer Protocol (TFTP) server, you can use either the **enable password** or **enable secret** global configuration commands. Both commands accomplish the same thing; that is, you can establish an encrypted password that users must enter to access privileged EXEC mode (the default) or any privilege level you specify.

We recommend that you use the **enable secret** command because it uses an improved encryption algorithm.

If you configure the **enable secret** command, it takes precedence over the **enable password** command; the two commands cannot be in effect simultaneously.

Use the **level** keyword to define a password for a specific privilege level. After you specify the level and set a password, give the password only to users who need to have access at this level. Use the **privilege level** global configuration command to specify commands accessible at various levels.

If you enable password encryption, it applies to all passwords including username passwords, authentication key passwords, the privileged command password, and console and virtual terminal line passwords.

To remove a password and level, use the **no enable password** [**level** *level*] or **no enable secret** [**level** *level*] global configuration command. To disable password encryption, use the **no service password-encryption** global configuration command.

Password Recovery

By default, any end user with physical access to the switch can recover from a lost password by interrupting the boot process while the switch is powering on and then by entering a new password.

The password-recovery disable feature protects access to the switch password by disabling part of this functionality. When this feature is enabled, the end user can interrupt the boot process only by agreeing to set the system back to the default configuration. With password recovery disabled, you can still interrupt the boot process and change the password, but the configuration file (config.text) and the VLAN database file (vlan.dat) are deleted.

Note If you disable password recovery, we recommend that you keep a backup copy of the configuration file on a secure server in case the end user interrupts the boot process and sets the system back to default values. Do not keep a backup copy of the configuration file on the switch. If the switch is operating in VTP transparent mode, we recommend that you also keep a backup copy of the VLAN database file on a secure server. When the switch is returned to the default system configuration, you can download the saved files to the switch by using the Xmodem protocol. For more information, see the "Recovering from a Lost or Forgotten Password" section on page 46-8.

Note Disabling password recovery will not work if you have set the switch to boot up manually by using the **boot manual** global configuration command. This command produces the boot loader prompt (*switch:*) after the switch is power cycled.

Telnet Password for a Terminal Line

When you power-up your switch for the first time, an automatic setup program runs to assign IP information and to create a default configuration for continued use. The setup program also prompts you to configure your switch for Telnet access through a password. If you did not configure this password during the setup program, you can configure it now through the command-line interface (CLI).

Username and Password Pairs

You can configure username and password pairs, which are locally stored on the switch. These pairs are assigned to lines or ports and authenticate each user before that user can access the switch. If you have defined privilege levels, you can also assign a specific privilege level (with associated rights and privileges) to each username and password pair.

Multiple Privilege Levels

By default, the Cisco IOS software has two modes of password security: user EXEC and privileged EXEC. You can configure up to 16 hierarchical levels of commands for each mode. By configuring multiple passwords, you can allow different sets of users to have access to specified commands.

For example, if you want many users to have access to the **clear line** command, you can assign it level 2 security and distribute the level 2 password fairly widely. But if you want more restricted access to the **configure** command, you can assign it level 3 security and distribute that password to a more restricted group of users.

When you set a command to a privilege level, all commands whose syntax is a subset of that command are also set to that level. For example, if you set the **show ip traffic** command to level 15, the **show** commands and **show ip** commands are automatically set to privilege level 15 unless you set them individually to different levels.

To return to the default privilege for a given command, use the **no privilege** *mode* **level** *level command* global configuration command.

Users can override the privilege level you set using the **privilege level** line configuration command by logging in to the line and enabling a different privilege level. They can lower the privilege level by using the **disable** command. If users know the password to a higher privilege level, they can use that password to enable the higher privilege level. You might specify a high level or privilege level for your console line to restrict line usage.

To return to the default line privilege level, use the **no privilege level** line configuration command.

Switch Access with TACACS+

This section describes how to enable and configure Terminal Access Controller Access Control System Plus (TACACS+), which provides detailed accounting information and flexible administrative control over authentication and authorization processes. TACACS+ is facilitated through authentication, authorization, accounting (AAA) and can be enabled only through AAA commands.

TACACS+

TACACS+ is a security application that provides centralized validation of users attempting to gain access to your switch. TACACS+ services are maintained in a database on a TACACS+ daemon typically running on a UNIX or Windows NT workstation. You should have access to and should configure a TACACS+ server before the configuring TACACS+ features on your switch.

TACACS+ provides for separate and modular authentication, authorization, and accounting facilities. TACACS+ allows for a single access control server (the TACACS+ daemon) to provide each service—authentication, authorization, and accounting—independently. Each service can be tied into its own database to take advantage of other services available on that server or on the network, depending on the capabilities of the daemon.

The goal of TACACS+ is to provide a method for managing multiple network access points from a single management service. Your switch can be a network access server along with other Cisco routers and access servers. A network access server provides connections to a single user, to a network or subnetwork, and to interconnected networks as shown in Figure 12-1.

Figure 12-1 Typical TACACS+ Network Configuration

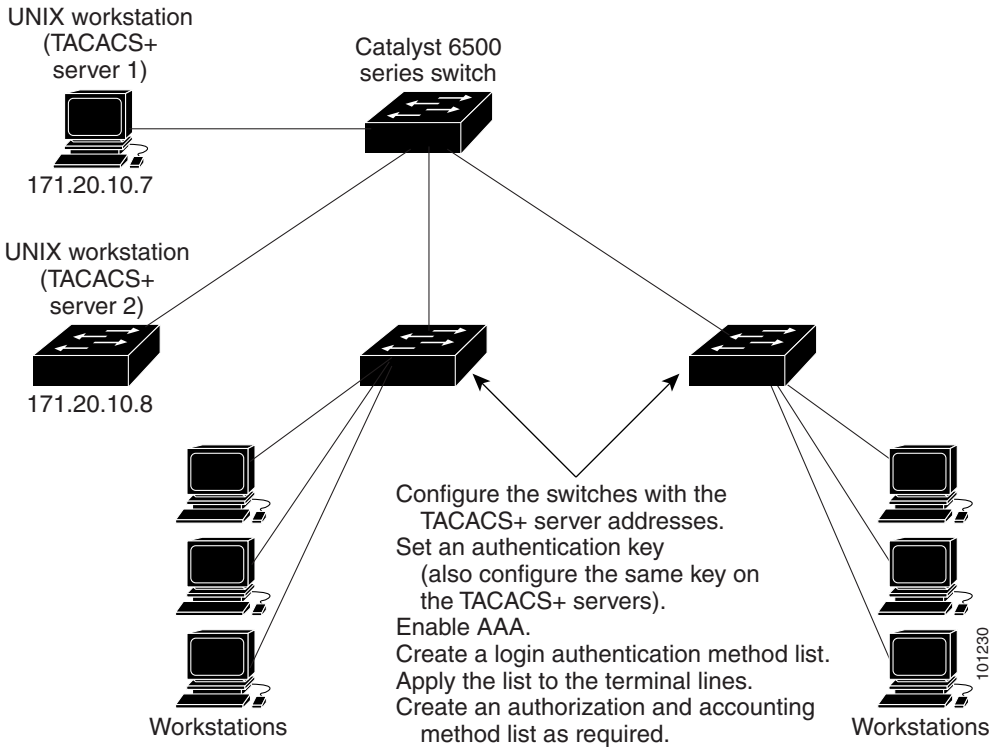

TACACS+, administered through the AAA security services, can provide these services:

- Authentication—Provides complete control of authentication through login and password dialog, challenge and response, and messaging support.

  The authentication facility can conduct a dialog with the user (for example, after a username and password are provided, to challenge a user with several questions, such as home address, mother's maiden name, service type, and social security number). The TACACS+ authentication service can also send messages to user screens. For example, a message could notify users that their passwords must be changed because of the company's password aging policy.

- Authorization—Provides fine-grained control over user capabilities for the duration of the user's session, including but not limited to setting autocommands, access control, session duration, or protocol support. You can also enforce restrictions on what commands a user can execute with the TACACS+ authorization feature.

- Accounting—Collects and sends information used for billing, auditing, and reporting to the TACACS+ daemon. Network managers can use the accounting facility to track user activity for a security audit or to provide information for user billing. Accounting records include user identities, start and stop times, executed commands (such as PPP), number of packets, and number of bytes.

The TACACS+ protocol provides authentication between the switch and the TACACS+ daemon, and it ensures confidentiality because all protocol exchanges between the switch and the TACACS+ daemon are encrypted.

You need a system running the TACACS+ daemon software to use TACACS+ on your switch.

TACACS+ Operation

When a user attempts a simple ASCII login by authenticating to a switch using TACACS+, this process occurs:

1. When the connection is established, the switch contacts the TACACS+ daemon to obtain a username prompt to show to the user. The user enters a username, and the switch then contacts the TACACS+ daemon to obtain a password prompt. The switch displays the password prompt to the user, the user enters a password, and the password is then sent to the TACACS+ daemon.

   TACACS+ allows a dialog between the daemon and the user until the daemon receives enough information to authenticate the user. The daemon prompts for a username and password combination, but can include other items, such as the user's mother's maiden name.

2. The switch eventually receives one of these responses from the TACACS+ daemon:

   - ACCEPT—The user is authenticated and service can begin. If the switch is configured to require authorization, authorization begins at this time.
   - REJECT—The user is not authenticated. The user can be denied access or is prompted to retry the login sequence, depending on the TACACS+ daemon.
   - ERROR—An error occurred at some time during authentication with the daemon or in the network connection between the daemon and the switch. If an ERROR response is received, the switch typically tries to use an alternative method for authenticating the user.
   - CONTINUE—The user is prompted for additional authentication information.

   After authentication, the user undergoes an additional authorization phase if authorization has been enabled on the switch. Users must first successfully complete TACACS+ authentication before proceeding to TACACS+ authorization.

3. If TACACS+ authorization is required, the TACACS+ daemon is again contacted, and it returns an ACCEPT or REJECT authorization response. If an ACCEPT response is returned, the response contains data in the form of attributes that direct the EXEC or NETWORK session for that user and the services that the user can access:
   - Telnet, Secure Shell (SSH), rlogin, or privileged EXEC services
   - Connection parameters, including the host or client IP address, access list, and user timeouts

Default TACACS+ Configuration

TACACS+ and AAA are disabled by default.

To prevent a lapse in security, you cannot configure TACACS+ through a network management application. When enabled, TACACS+ can authenticate users accessing the switch through the CLI.

Note Although TACACS+ configuration is performed through the CLI, the TACACS+ server authenticates HTTP connections that have been configured with a privilege level of 15.

TACACS+ Server Host and the Authentication Key

You can configure the switch to use a single server or AAA server groups to group existing server hosts for authentication. You can group servers to select a subset of the configured server hosts and use them for a particular service. The server group is used with a global server-host list and contains the list of IP addresses of the selected server hosts.

TACACS+ Login Authentication

To configure AAA authentication, you define a named list of authentication methods and then apply that list to various ports. The method list defines the types of authentication to be performed and the sequence in which they are performed; it must be applied to a specific port before any of the defined authentication methods are performed. The only exception is the default method list (which, by coincidence, is named *default*). The default method list is automatically applied to all ports except those that have a named method list explicitly defined. A defined method list overrides the default method list.

A method list describes the sequence and authentication methods to be queried to authenticate a user. You can designate one or more security protocols to be used for authentication, thus ensuring a backup system for authentication in case the initial method fails. The software uses the first method listed to authenticate users; if that method fails to respond, the software selects the next authentication method in the method list. This process continues until there is successful communication with a listed authentication method or until all defined methods are exhausted. If authentication fails at any point in this cycle—meaning that the security server or local username database responds by denying the user access—the authentication process stops, and no other authentication methods are attempted.

TACACS+ Authorization for Privileged EXEC Access and Network Services

AAA authorization limits the services available to a user. When AAA authorization is enabled, the switch uses information retrieved from the user's profile, which is located either in the local user database or on the security server, to configure the user's session. The user is granted access to a requested service only if the information in the user profile allows it.

You can use the **aaa authorization** global configuration command with the **tacacs+** keyword to set parameters that restrict a user's network access to privileged EXEC mode.

The **aaa authorization exec tacacs+ local** command sets these authorization parameters:

- Use TACACS+ for privileged EXEC access authorization if authentication was performed by using TACACS+.
- Use the local database if authentication was not performed by using TACACS+.

Note Authorization is bypassed for authenticated users who log in through the CLI even if authorization has been configured.

TACACS+ Accounting

The AAA accounting feature tracks the services that users are accessing and the amount of network resources that they are consuming. When AAA accounting is enabled, the switch reports user activity to the TACACS+ security server in the form of accounting records. Each accounting record contains accounting attribute-value (AV) pairs and is stored on the security server. This data can then be analyzed for network management, client billing, or auditing.

Switch Access with RADIUS

This section describes how to enable and configure the RADIUS, which provides detailed accounting information and flexible administrative control over authentication and authorization processes. RADIUS is facilitated through AAA and can be enabled only through AAA commands.

RADIUS

RADIUS is a distributed client/server system that secures networks against unauthorized access. RADIUS clients run on supported Cisco routers and switches. Clients send authentication requests to a central RADIUS server, which contains all user authentication and network service access information. The RADIUS host is normally a multiuser system running RADIUS server software from Cisco (Cisco Secure Access Control Server Version 3.0), Livingston, Merit, Microsoft, or another software provider. For more information, see the RADIUS server documentation.

Use RADIUS in these network environments that require access security:

- Networks with multiple-vendor access servers, each supporting RADIUS. For example, access servers from several vendors use a single RADIUS server-based security database. In an IP-based network with multiple vendors' access servers, dial-in users are authenticated through a RADIUS server that has been customized to work with the Kerberos security system.
- Turnkey network security environments in which applications support the RADIUS protocol, such as in an access environment that uses a *smart card* access control system. In one case, RADIUS has been used with Enigma's security cards to validates users and to grant access to network resources.
- Networks already using RADIUS. You can add a Cisco switch containing a RADIUS client to the network. This might be the first step when you make a transition to a TACACS+ server.
- Network in which the user must only access a single service. Using RADIUS, you can control user access to a single host, to a single utility such as Telnet, or to the network through a protocol such as IEEE 802.1x. For more information about this protocol, see Chapter 13, "Configuring IEEE 802.1x Port-Based Authentication."

- Networks that require resource accounting. You can use RADIUS accounting independently of RADIUS authentication or authorization. The RADIUS accounting functions allow data to be sent at the start and end of services, showing the amount of resources (such as time, packets, bytes, and so forth) used during the session. An Internet service provider might use a freeware-based version of RADIUS access control and accounting software to meet special security and billing needs.

RADIUS is not suitable in these network security situations:

- Multiprotocol access environments. RADIUS does not support AppleTalk Remote Access (ARA), NetBIOS Frame Control Protocol (NBFCP), NetWare Asynchronous Services Interface (NASI), or X.25 PAD connections.
- Switch-to-switch or router-to-router situations. RADIUS does not provide two-way authentication. RADIUS can be used to authenticate from one device to a non-Cisco device if the non-Cisco device requires authentication.
- Networks using a variety of services. RADIUS generally binds a user to one service model.

Figure 12-2 Transitioning from RADIUS to TACACS+ Services

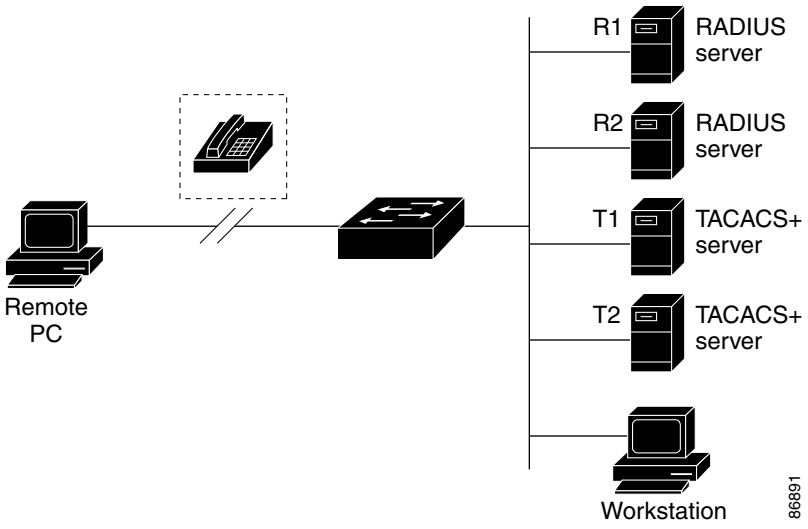

RADIUS Operation

When a user attempts to log in and authenticate to a switch that is access controlled by a RADIUS server, these events occur:

1. The user is prompted to enter a username and password.
2. The username and encrypted password are sent over the network to the RADIUS server.
3. The user receives one of these responses from the RADIUS server:
   a. ACCEPT—The user is authenticated.
   b. REJECT—The user is either not authenticated and is prompted to re-enter the username and password, or access is denied.
   c. CHALLENGE—A challenge requires additional data from the user.
   d. CHALLENGE PASSWORD—A response requests the user to select a new password.

The ACCEPT or REJECT response is bundled with additional data that is used for privileged EXEC or network authorization. Users must first successfully complete RADIUS authentication before proceeding to RADIUS authorization, if it is enabled. The additional data included with the ACCEPT or REJECT packets includes these items:

- Telnet, SSH, rlogin, or privileged EXEC services
- Connection parameters, including the host or client IP address, access list, and user timeouts

Default RADIUS Configuration

RADIUS and AAA are disabled by default.

To prevent a lapse in security, you cannot configure RADIUS through a network management application. When enabled, RADIUS can authenticate users accessing the switch through the CLI.

RADIUS Change of Authorization

This section provides an overview of the RADIUS interface including available primitives and how they are used during a Change of Authorization (CoA).

Radius COA Overview

A standard RADIUS interface is typically used in a pulled model where the request originates from a network attached device and the response come from the queried servers. Catalyst switches support the RADIUS Change of Authorization (CoA) extensions defined in RFC 5176 that are typically used in a pushed model and allow for the dynamic reconfiguring of sessions from external authentication, authorization, and accounting (AAA) or policy servers.

The switch supports these per-session CoA requests:

- Session reauthentication
- Session termination
- Session termination with port shutdown
- Session termination with port bounce

Change-of-Authorization Requests

Change of Authorization (CoA) requests, as described in RFC 5176, are used in a push model to allow for session identification, host reauthentication, and session termination. The model is comprised of one request (CoA-Request) and two possible response codes:

- CoA acknowledgement (ACK) [CoA-ACK]
- CoA non-acknowledgement (NAK) [CoA-NAK]

The request is initiated from a CoA client (typically a RADIUS or policy server) and directed to the switch that acts as a listener.

RFC 5176 Compliance

The Disconnect Request message, which is also referred to as Packet of Disconnect (POD), is supported by the switch for session termination.

Table 12-2 Supported IETF Attributes

Attribute Number	Attribute Name
24	State
31	Calling-Station-ID
44	Acct-Session-ID
80	Message-Authenticator
101	Error-Cause

Table 12-3 Error-Cause Values

Value	Explanation
201	Residual Session Context Removed
202	Invalid EAP Packet (Ignored)
401	Unsupported Attribute
402	Missing Attribute
403	NAS Identification Mismatch
404	Invalid Request
405	Unsupported Service
406	Unsupported Extension
407	Invalid Attribute Value
501	Administratively Prohibited
502	Request Not Routable (Proxy)
503	Session Context Not Found
504	Session Context Not Removable
505	Other Proxy Processing Error
506	Resources Unavailable
507	Request Initiated
508	Multiple Session Selection Unsupported

CoA Request Response Code

The CoA Request response code can be used to convey a command to the switch. The supported commands are listed in Table 12-4 on page 12-12.

CoA Session Identification

For disconnect and CoA requests targeted at a particular session, the switch locates the session based on one or more of the following attributes:

- Calling-Station-Id (IETF attribute 31 which contains the host MAC address)
- Audit-Session-Id (Cisco VSA)
- Acct-Session-Id (IETF attribute 44)

Unless all session identification attributes included in the CoA message match the session, the switch returns a Disconnect-NAK or CoA-NAK with the Invalid Attribute Value error-code attribute.

For disconnect and CoA requests targeted to a particular session, any one of these session identifiers can be used:

- Calling-Station-ID (IETF attribute 31, which should contain the MAC address)
- Audit-Session-ID (Cisco vendor-specific attribute)
- Accounting-Session-ID (IETF attribute 44).

If more than one session identification attribute is included in the message, all the attributes must match the session or the switch returns a Disconnect- negative acknowledgement (NAK) or CoA-NAK with the error code *Invalid Attribute Value*.

The packet format for a CoA Request code as defined in RFC 5176 consists of the fields: Code, Identifier, Length, Authenticator, and Attributes in Type:Length:Value (TLV) format.

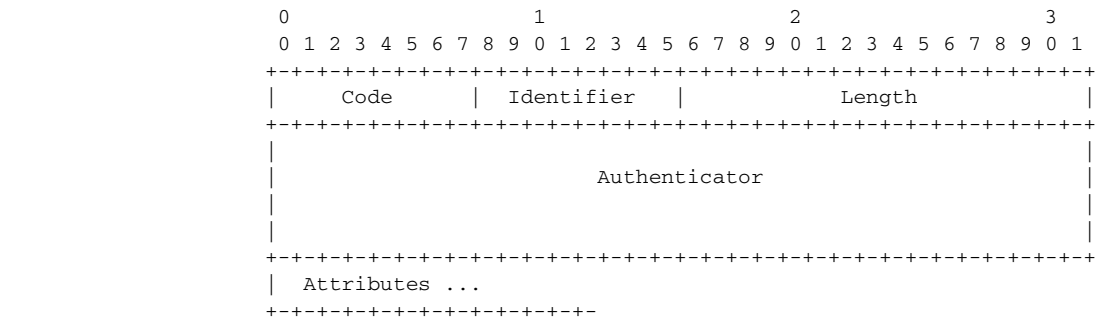

The attributes field is used to carry Cisco VSAs.

CoA ACK Response Code

If the authorization state is changed successfully, a positive acknowledgement (ACK) is sent. The attributes returned within CoA ACK will vary based on the CoA Request and are discussed in individual CoA Commands.

CoA NAK Response Code

A negative acknowledgement (NAK) indicates a failure to change the authorization state and can include attributes that indicate the reason for the failure. Use **show** commands to verify a successful CoA.

CoA Request Commands

Table 12-4 CoA Commands Supported on the Switch

Command[1]	Cisco VSA
Reauthenticate host	Cisco:Avpair="subscriber:command=reauthenticate"
Terminate session	This is a standard disconnect request that does not require a VSA.
Bounce host port	Cisco:Avpair="subscriber:command=bounce-host-port"
Disable host port	Cisco:Avpair="subscriber:command=disable-host-port"

1. All CoA commands must include the session identifier between the switch and the CoA client.

CoA Session Reauthentication

The AAA server typically generates a session reauthentication request when a host with an unknown identity or posture joins the network and is associated with a restricted access authorization profile (such as a guest VLAN). A reauthentication request allows the host to be placed in the appropriate authorization group when its credentials are known.

To initiate session authentication, the AAA server sends a standard CoA-Request message which contains a Cisco vendor-specific attribute (VSA) in this form:
Cisco:Avpair="subscriber:command=reauthenticate" and one or more session identification attributes.

The current session state determines the switch response to the message. If the session is currently authenticated by IEEE 802.1x, the switch responds by sending an Extensible Authentication Protocol over LAN (EAPoL) RequestId message to the server.

If the session is currently authenticated by MAC authentication bypass (MAB), the switch sends an access-request to the server, passing the same identity attributes used for the initial successful authentication.

If session authentication is in progress when the switch receives the command, the switch terminates the process, and restarts the authentication sequence, starting with the method configured to be attempted first.

If the session is not yet authorized, or is authorized via guest VLAN, or critical VLAN, or similar policies, the reauthentication message restarts the access control methods, beginning with the method configured to be attempted first. The current authorization of the session is maintained until the reauthentication leads to a different authorization result.

CoA Session Termination

There are three types of CoA requests that can trigger session termination. A CoA Disconnect-Request terminates the session, without disabling the host port. This command causes reinitialization of the authenticator state machine for the specified host, but does not restrict that host's access to the network.

To restrict a host's access to the network, use a CoA Request with the Cisco:Avpair="subscriber:command=disable-host-port" VSA. This command is useful when a host is known to be causing problems on the network, and you need to immediately block network access for the host. When you want to restore network access on the port, reenable it using a non-RADIUS mechanism.

When a device with no supplicant, such as a printer, needs to acquire a new IP address (for example, after a VLAN change), terminate the session on the host port with port-bounce (temporarily disable and then reenable the port).

CoA Disconnect-Request

This command is a standard Disconnect-Request. Because this command is session-oriented, it must be accompanied by one or more of the session identification attributes described in the "CoA Session Identification" section on page 12-11. If the session cannot be located, the switch returns a Disconnect-NAK message with the "Session Context Not Found" error-code attribute. If the session *is* located, the switch terminates the session. After the session has been completely removed, the switch returns a Disconnect-ACK.

If the switch fails-over to a standby switch before returning a Disconnect-ACK to the client, the process is repeated on the new active switch when the request is resent from the client. If the session is not found following resend, a Disconnect-ACK is sent with the "Session Context Not Found" error-code attribute.

CoA Request: Disable Host Port

This command is carried in a standard CoA-Request message that has this new VSA:

Cisco:Avpair="subscriber:command=disable-host-port"

Because this command is session-oriented, it must be accompanied by one or more of the session identification attributes described in the "CoA Session Identification" section on page 12-11. If the session cannot be located, the switch returns a CoA-NAK message with the "Session Context Not Found" error-code attribute. If the session is located, the switch disables the hosting port and returns a CoA-ACK message.

If the switch fails before returning a CoA-ACK to the client, the process is repeated on the new active switch when the request is resent from the client. If the switch fails after returning a CoA-ACK message to the client but before the operation has completed, the operation is restarted on the new active switch.

Note A Disconnect-Request failure following command resend could be the result of either a successful session termination before change-over (if the Disconnect-ACK was not sent) or a session termination by other means (for example, a link failure) that occurred after the original command was issued and before the standby switch became active.

CoA Request: Bounce-Port

This command is carried in a standard CoA-Request message that contains this VSA:
Cisco:Avpair="subscriber:command=bounce-host-port"

Because this command is session-oriented, it must be accompanied by one or more of the session identification attributes described in the "CoA Session Identification" section on page 12-11. If the session cannot be located, the switch returns a CoA-NAK message with the "Session Context Not Found" error-code attribute. If the session is located, the switch disables the hosting port for a period of 10 seconds, reenables it (port-bounce), and returns a CoA-ACK.

If the switch fails before returning a CoA-ACK to the client, the process is repeated on the new active switch when the request is resent from the client. If the switch fails after returning a CoA-ACK message to the client but before the operation has completed, the operation is restarted on the new active switch.

RADIUS Server Host

Switch-to-RADIUS-server communication involves several components:

- Hostname or IP address
- Authentication destination port
- Accounting destination port
- Key string
- Timeout period
- Retransmission value

You identify RADIUS security servers by their hostname or IP address, hostname and specific UDP port numbers, or their IP address and specific UDP port numbers. The combination of the IP address and the UDP port number creates a unique identifier, allowing different ports to be individually defined as RADIUS hosts providing a specific AAA service. This unique identifier enables RADIUS requests to be sent to multiple UDP ports on a server at the same IP address.

If two different host entries on the same RADIUS server are configured for the same service—for example, accounting—the second host entry configured acts as a fail-over backup to the first one. Using this example, if the first host entry fails to provide accounting services, the `%RADIUS-4-RADIUS_DEAD` message appears, and then the switch tries the second host entry configured on the same device for accounting services. (The RADIUS host entries are tried in the order that they are configured.)

A RADIUS server and the switch use a shared secret text string to encrypt passwords and exchange responses. To configure RADIUS to use the AAA security commands, you must specify the host running the RADIUS server daemon and a secret text (key) string that it shares with the switch.

The timeout, retransmission, and encryption key values can be configured globally for all RADIUS servers, on a per-server basis, or in some combination of global and per-server settings. To apply these settings globally to all RADIUS servers communicating with the switch, use the three unique global configuration commands: **radius-server timeout**, **radius-server retransmit**, and **radius-server key**. To apply these values on a specific RADIUS server, use the **radius-server host** global configuration command.

Note If you configure both global and per-server functions (timeout, retransmission, and key commands) on the switch, the per-server timer, retransmission, and key value commands override global timer, retransmission, and key value commands. For information on configuring these settings on all RADIUS servers, see the "Configuring Settings for All RADIUS Servers" section on page 12-37.

You can configure the switch to use AAA server groups to group existing server hosts for authentication. For more information, see the "Defining AAA Server Groups" section on page 12-35.

RADIUS Login Authentication

To configure AAA authentication, you define a named list of authentication methods and then apply that list to various ports. The method list defines the types of authentication to be performed and the sequence in which they are performed; it must be applied to a specific port before any of the defined authentication methods are performed. The only exception is the default method list (which, by coincidence, is named *default*). The default method list is automatically applied to all ports except those that have a named method list explicitly defined.

Radius Method List

A method list defines the sequence and methods to be used to authenticate, to authorize, or to keep accounts on a user. You can use method lists to designate one or more security protocols to be used (such as TACACS+ or local username lookup), which ensures a backup system if the initial method fails. The software uses the first method listed to authenticate, to authorize, or to keep accounts on users. If that method does not respond, the software selects the next method in the list. This process continues until there is successful communication with a listed method or the method list is exhausted.

AAA Server Groups

You can configure the switch to use AAA server groups to group existing server hosts for authentication. You select a subset of the configured server hosts and use them for a particular service. The server group is used with a global server-host list, which lists the IP addresses of the selected server hosts.

Server groups also can include multiple host entries for the same server if each entry has a unique identifier (the combination of the IP address and UDP port number), allowing different ports to be individually defined as RADIUS hosts providing a specific AAA service. If you configure two different host entries on the same RADIUS server for the same service, (for example, accounting), the second configured host entry acts as a failover backup to the first one.

You use the **server** group server configuration command to associate a particular server with a defined group server. You can either identify the server by its IP address or identify multiple host instances or entries by using the optional **auth-port** and **acct-port** keywords.

RADIUS Authorization for User Privileged Access and Network Services

AAA authorization limits the services available to a user. When AAA authorization is enabled, the switch uses information retrieved from the user's profile, which is in the local user database or on the security server, to configure the user's session. The user is granted access to a requested service only if the information in the user profile allows it.

You can use the **aaa authorization** global configuration command with the **radius** keyword to set parameters that restrict a user's network access to privileged EXEC mode.

The **aaa authorization exec radius local** command sets these authorization parameters:

- Use RADIUS for privileged EXEC access authorization if authentication was performed by using RADIUS.
- Use the local database if authentication was not performed by using RADIUS.

Note Authorization is bypassed for authenticated users who log in through the CLI even if authorization has been configured.

RADIUS Accounting

The AAA accounting feature tracks the services that users are accessing and the amount of network resources that they are consuming. When AAA accounting is enabled, the switch reports user activity to the RADIUS security server in the form of accounting records. Each accounting record contains accounting attribute-value (AV) pairs and is stored on the security server. This data can then be analyzed for network management, client billing, or auditing.

Establishing a Session with a Router if the AAA Server is Unreachable

The **aaa accounting system guarantee-first** command guarantees system accounting as the first record, which is the default condition. In some situations, users might be prevented from starting a session on the console or terminal connection until after the system reloads, which can take more than 3 minutes.

To establish a console or Telnet session with the router if the AAA server is unreachable when the router reloads, use the **no aaa accounting system guarantee-first** command.

Vendor-Specific RADIUS Attributes

The Internet Engineering Task Force (IETF) draft standard specifies a method for communicating vendor-specific information between the switch and the RADIUS server by using the vendor-specific attribute (attribute 26). Vendor-specific attributes (VSAs) allow vendors to support their own extended

attributes not suitable for general use. The Cisco RADIUS implementation supports one vendor-specific option by using the format recommended in the specification. Cisco's vendor-ID is 9, and the supported option has vendor-type 1, which is named *cisco-avpair*. The value is a string with this format:

```
protocol : attribute sep value *
```

protocol is a value of the Cisco protocol attribute for a particular type of authorization. *Attribute* and *value* are an appropriate attribute-value (AV) pair defined in the Cisco TACACS+ specification, and *sep* is = for mandatory attributes and is * for optional attributes. The full set of features available for TACACS+ authorization can then be used for RADIUS.

For example, this AV pair activates Cisco's *multiple named ip address pools* feature during IP authorization (during PPP IPCP address assignment):

```
cisco-avpair= "ip:addr-pool=first"
```

Vendor-Proprietary RADIUS Server Communication

Although an IETF draft standard for RADIUS specifies a method for communicating vendor-proprietary information between the switch and the RADIUS server, some vendors have extended the RADIUS attribute set in a unique way. Cisco IOS software supports a subset of vendor-proprietary RADIUS attributes.

As mentioned earlier, to configure RADIUS (whether vendor-proprietary or IETF draft-compliant), you must specify the host running the RADIUS server daemon and the secret text string it shares with the switch. You specify the RADIUS host and secret text string by using the **radius-server** global configuration commands.

Switch Access with Kerberos

This section describes how to enable and configure the Kerberos security system, which authenticates requests for network resources by using a trusted third party. To use this feature, the cryptographic (that is, supports encryption) versions of the switch software must be installed on your switch.

You must obtain authorization to use this feature and to download the cryptographic software files from Cisco.com. For more information, see the release notes for this release.

Understanding Kerberos

Kerberos is a secret-key network authentication protocol, which was developed at the Massachusetts Institute of Technology (MIT). It uses the Data Encryption Standard (DES) cryptographic algorithm for encryption and authentication and authenticates requests for network resources. Kerberos uses the concept of a trusted third party to perform secure verification of users and services. This trusted third party is called the *key distribution center* (KDC).

Kerberos verifies that users are who they claim to be and the network services that they use are what the services claim to be. To do this, a KDC or trusted Kerberos server issues tickets to users. These tickets, which have a limited lifespan, are stored in user credential caches. The Kerberos server uses the tickets instead of usernames and passwords to authenticate users and network services.

Note A Kerberos server can be a switch that is configured as a network security server and that can authenticate users by using the Kerberos protocol.

The Kerberos credential scheme uses a process called *single logon*. This process authenticates a user once and then allows secure authentication (without encrypting another password) wherever that user credential is accepted.

This software release supports Kerberos 5, which allows organizations that are already using Kerberos 5 to use the same Kerberos authentication database on the KDC that they are already using on their other network hosts (such as UNIX servers and PCs).

In this software release, Kerberos supports these network services:

- Telnet
- rlogin
- rsh (Remote Shell Protocol)

Table 12-5 lists the common Kerberos-related terms and definitions.

Table 12-5 Kerberos Terms

Term	Definition
Authentication	A process by which a user or service identifies itself to another service. For example, a client can authenticate to a switch or a switch can authenticate to another switch.
Authorization	A means by which the switch identifies what privileges the user has in a network or on the switch and what actions the user can perform.
Credential	A general term that refers to authentication tickets, such as TGTs[1] and service credentials. Kerberos credentials verify the identity of a user or service. If a network service decides to trust the Kerberos server that issued a ticket, it can be used in place of reentering a username and password. Credentials have a default lifespan of eight hours.
Instance	An authorization level label for Kerberos principals. Most Kerberos principals are of the form *user@REALM* (for example, smith@EXAMPLE.COM). A Kerberos principal with a Kerberos instance has the form *user/instance@REALM* (for example, smith/admin@EXAMPLE.COM). The Kerberos instance can be used to specify the authorization level for the user if authentication is successful. The server of each network service might implement and enforce the authorization mappings of Kerberos instances but is not required to do so. **Note** The Kerberos principal and instance names *must* be in all lowercase characters. **Note** The Kerberos realm name *must* be in all uppercase characters.
KDC[2]	Key distribution center that consists of a Kerberos server and database program that is running on a network host.
Kerberized	A term that describes applications and services that have been modified to support the Kerberos credential infrastructure.
Kerberos realm	A domain consisting of users, hosts, and network services that are registered to a Kerberos server. The Kerberos server is trusted to verify the identity of a user or network service to another user or network service. **Note** The Kerberos realm name *must* be in all uppercase characters.

Table 12-5 Kerberos Terms (continued)

Term	Definition
Kerberos server	A daemon that is running on a network host. Users and network services register their identity with the Kerberos server. Network services query the Kerberos server to authenticate to other network services.
KEYTAB[3]	A password that a network service shares with the KDC. In Kerberos 5 and later Kerberos versions, the network service authenticates an encrypted service credential by using the KEYTAB to decrypt it. In Kerberos versions earlier than Kerberos 5, KEYTAB is referred to as SRVTAB[4].
Principal	Also known as a Kerberos identity, this is who you are or what a service is according to the Kerberos server. **Note** The Kerberos principal name *must* be in all lowercase characters.
Service credential	A credential for a network service. When issued from the KDC, this credential is encrypted with the password shared by the network service and the KDC. The password is also shared with the user TGT.
SRVTAB	A password that a network service shares with the KDC. In Kerberos 5 or later Kerberos versions, SRVTAB is referred to as KEYTAB.
TGT	Ticket granting ticket that is a credential that the KDC issues to authenticated users. When users receive a TGT, they can authenticate to network services within the Kerberos realm represented by the KDC.

1. TGT = ticket granting ticket
2. KDC = key distribution center
3. KEYTAB = key table
4. SRVTAB = server table

Kerberos Operation

A Kerberos server can be a switch that is configured as a network security server and that can authenticate remote users by using the Kerberos protocol. Although you can customize Kerberos in a number of ways, remote users attempting to access network services must pass through three layers of security before they can access network services.

To authenticate to network services by using a switch as a Kerberos server, remote users must follow these steps:

1. Authenticating to a Boundary Switch, page 12-19
2. Obtaining a TGT from a KDC, page 12-20
3. Authenticating to Network Services, page 12-20

Authenticating to a Boundary Switch

This section describes the first layer of security through which a remote user must pass. The user must first authenticate to the boundary switch. This process then occurs:

1. The user opens an un-Kerberized Telnet connection to the boundary switch.
2. The switch prompts the user for a username and password.
3. The switch requests a TGT from the KDC for this user.

4. The KDC sends an encrypted TGT that includes the user identity to the switch.
5. The switch attempts to decrypt the TGT by using the password that the user entered.
   - If the decryption is successful, the user is authenticated to the switch.
   - If the decryption is not successful, the user repeats Step 2 either by reentering the username and password (noting if Caps Lock or Num Lock is on or off) or by entering a different username and password.

A remote user who initiates a un-Kerberized Telnet session and authenticates to a boundary switch is inside the firewall, but the user must still authenticate directly to the KDC before getting access to the network services. The user must authenticate to the KDC because the TGT that the KDC issues is stored on the switch and cannot be used for additional authentication until the user logs on to the switch.

Obtaining a TGT from a KDC

This section describes the second layer of security through which a remote user must pass. The user must now authenticate to a KDC and obtain a TGT from the KDC to access network services.

Authenticating to Network Services

This section describes the third layer of security through which a remote user must pass. The user with a TGT must now authenticate to the network services in a Kerberos realm.

Kerberos Configuration

So that remote users can authenticate to network services, you must configure the hosts and the KDC in the Kerberos realm to communicate and mutually authenticate users and network services. To do this, you must identify them to each other. You add entries for the hosts to the Kerberos database on the KDC and add KEYTAB files generated by the KDC to all hosts in the Kerberos realm. You also create entries for the users in the KDC database.

When you add or create entries for the hosts and users, follow these guidelines:

- The Kerberos principal name *must* be in all lowercase characters.
- The Kerberos instance name *must* be in all lowercase characters.
- The Kerberos realm name *must* be in all uppercase characters.

Note A Kerberos server can be a switch that is configured as a network security server and that can authenticate users by using the Kerberos protocol.

To set up a Kerberos-authenticated server-client system, follow these steps:

- Configure the KDC by using Kerberos commands.
- Configure the switch to use the Kerberos protocol.

Local Authentication and Authorization

You can configure AAA to operate without a server by setting the switch to implement AAA in local mode. The switch then handles authentication and authorization. No accounting is available in this configuration.

Secure Shell

To use this feature, you must install the cryptographic (encrypted) software image on your switch. You must obtain authorization to use this feature and to download the cryptographic software files from Cisco.com. For more information, see the release notes for this release.

For SSH configuration examples, see the "SSH Configuration Examples" section in the "Configuring Secure Shell" chapter of the *Cisco IOS Security Configuration Guide, Cisco IOS Release 12.2*.

SSH in IPv6 functions the same and offers the same benefits as SSH in IPv4. IPv6 enhancements to SSH consist of support for IPv6 addresses that enable a Cisco router to accept and establish secure, encrypted connections with remote IPv6 nodes over an IPv6 transport.

Note For complete syntax and usage information for the commands used in this section, see the command reference for this release and command reference for Cisco IOS Release 12.2.

SSH

SSH is a protocol that provides a secure, remote connection to a device. SSH provides more security for remote connections than Telnet does by providing strong encryption when a device is authenticated. This software release supports SSH Version 1 (SSHv1) and SSH Version 2 (SSHv2).

SSH Servers, Integrated Clients, and Supported Versions

The SSH feature has an SSH server and an SSH integrated client, which are applications that run on the switch. You can use an SSH client to connect to a switch running the SSH server. The SSH server works with the SSH client supported in this release and with non-Cisco SSH clients. The SSH client also works with the SSH server supported in this release and with non-Cisco SSH servers.

The switch supports an SSHv1 or an SSHv2 server.

The switch supports an SSHv1 client.

SSH supports the Data Encryption Standard (DES) encryption algorithm, the Triple DES (3DES) encryption algorithm, and password-based user authentication.

SSH also supports these user authentication methods:

- TACACS+ (for more information, see the "Configuring TACACS+" section on page 12-30)
- RADIUS (for more information, see the "Configuring Radius Server Communication" section on page 12-33)
- Local authentication and authorization (for more information, see the "Configuring the Switch for Local Authentication and Authorization" section on page 12-39)

Note This software release does not support IP Security (IPSec).

Limitations

These limitations apply to SSH:

- The switch supports Rivest, Shamir, and Adelman (RSA) authentication.
- SSH supports only the execution-shell application.
- The SSH server and the SSH client are supported only on DES (56-bit) and 3DES (168-bit) data encryption software.
- The switch supports the Advanced Encryption Standard (AES) encryption algorithm with a 128-bit key, 192-bit key, or 256-bit key. However, symmetric cipher AES to encrypt the keys is not supported.

SSH Configuration Guidelines

Follow these guidelines when configuring the switch as an SSH server or SSH client:

- An RSA key pair generated by a SSHv1 server can be used by an SSHv2 server, and the reverse.
- If you get CLI error messages after entering the **crypto key generate rsa** global configuration command, an RSA key pair has not been generated. Reconfigure the hostname and domain, and then enter the **crypto key generate rsa** command. For more information, see the "Setting Up the Switch to Run SSH" section on page 12-40.
- When generating the RSA key pair, the message `No host name specified` might appear. If it does, you must configure a hostname by using the **hostname** global configuration command.
- When generating the RSA key pair, the message `No domain specified` might appear. If it does, you must configure an IP domain name by using the **ip domain-name** global configuration command.
- When configuring the local authentication and authorization authentication method, make sure that AAA is disabled on the console.

Switch for Secure Socket Layer HTTP

Secure Socket Layer (SSL) version 3.0 supports the HTTP 1.1 server and client. SSL provides server authentication, encryption, and message integrity, as well as HTTP client authentication, to allow secure HTTP communications. To use this feature, the cryptographic (encrypted) software image must be installed on your switch. You must obtain authorization to use this feature and to download the cryptographic software files from Cisco.com. For more information about the crypto image, see the release notes for this release.

Secure HTTP Servers and Clients

On a secure HTTP connection, data to and from an HTTP server is encrypted before being sent over the Internet. HTTP with SSL encryption provides a secure connection to allow such functions as configuring a switch from a Web browser. Cisco's implementation of the secure HTTP server and secure HTTP client uses an implementation of SSL Version 3.0 with application-layer encryption. HTTP over SSL is abbreviated as HTTPS; the URL of a secure connection begins with https:// instead of http://.

The primary role of the HTTP secure server (the switch) is to listen for HTTPS requests on a designated port (the default HTTPS port is 443) and pass the request to the HTTP 1.1 Web server. The HTTP 1.1 server processes requests and passes responses (pages) back to the HTTP secure server, which responds to the original request.

The primary role of the HTTP secure client (the web browser) is to respond to Cisco IOS application requests for HTTPS User Agent services, perform HTTPS User Agent services for the application, and pass the response back to the application.

When SSL is used in a switch cluster, the SSL session terminates at the cluster commander. Cluster member switches must run standard HTTP.

For secure HTTP connections, we recommend that you configure an official CA trustpoint. A CA trustpoint is more secure than a self-signed certificate.

Before you configure a CA trustpoint, you should ensure that the system clock is set. If the clock is not set, the certificate is rejected due to an incorrect date.

Default SSL Settings

Table 12-6 Default SSL Settings

Default Setting
The standard HTTP server is enabled.
SSL is enabled.
No CA trustpoints are configured.
No self-signed certificates are generated.

Certificate Authority Trustpoints

Certificate authorities (CAs) manage certificate requests and issue certificates to participating network devices. These services provide centralized security key and certificate management for the participating devices. Specific CA servers are referred to as *trustpoints*.

When a connection attempt is made, the HTTPS server provides a secure connection by issuing a certified X.509v3 certificate, obtained from a specified CA trustpoint, to the client. The client (usually a Web browser), in turn, has a public key that allows it to authenticate the certificate.

For secure HTTP connections, we highly recommend that you configure a CA trustpoint. If a CA trustpoint is not configured for the device running the HTTPS server, the server certifies itself and generates the needed RSA key pair. Because a self-certified (self-signed) certificate does not provide adequate security, the connecting client generates a notification that the certificate is self-certified, and the user has the opportunity to accept or reject the connection. This option is useful for internal network topologies (such as testing).

If you do not configure a CA trustpoint, when you enable a secure HTTP connection, either a temporary or a persistent self-signed certificate for the secure HTTP server (or client) is automatically generated.

- If the switch is not configured with a hostname and a domain name, a temporary self-signed certificate is generated. If the switch reboots, any temporary self-signed certificate is lost, and a new temporary new self-signed certificate is assigned.
- If the switch has been configured with a host and domain name, a persistent self-signed certificate is generated. This certificate remains active if you reboot the switch or if you disable the secure HTTP server so that it will be there the next time you reenable a secure HTTP connection.

Note The certificate authorities and trustpoints must be configured on each device individually. Copying them from other devices makes them invalid on the switch.

Note The values that follow *TP self-signed* depend on the serial number of the device.

You can use an optional command (**ip http secure-client-auth**) to allow the HTTPS server to request an X.509v3 certificate from the client. Authenticating the client provides more security than server authentication by itself.

CipherSuites

A CipherSuite specifies the encryption algorithm and the digest algorithm to use on a SSL connection. When connecting to the HTTPS server, the client Web browser offers a list of supported CipherSuites, and the client and server negotiate the best encryption algorithm to use from those on the list that are supported by both. For example, Netscape Communicator 4.76 supports U.S. security with RSA Public Key Cryptography, MD2, MD5, RC2-CBC, RC4, DES-CBC, and DES-EDE3-CBC.

For the best possible encryption, you should use a client browser that supports 128-bit encryption, such as Microsoft Internet Explorer Version 5.5 (or later) or Netscape Communicator Version 4.76 (or later). The SSL_RSA_WITH_DES_CBC_SHA CipherSuite provides less security than the other CipherSuites, as it does not offer 128-bit encryption.

The more secure and more complex CipherSuites require slightly more processing time. This list defines the CipherSuites supported by the switch and ranks them from fastest to slowest in terms of router processing load (speed):

1. SSL_RSA_WITH_DES_CBC_SHA—RSA key exchange (RSA Public Key Cryptography) with DES-CBC for message encryption and SHA for message digest
2. SSL_RSA_WITH_RC4_128_MD5—RSA key exchange with RC4 128-bit encryption and MD5 for message digest
3. SSL_RSA_WITH_RC4_128_SHA—RSA key exchange with RC4 128-bit encryption and SHA for message digest
4. SSL_RSA_WITH_3DES_EDE_CBC_SHA—RSA key exchange with 3DES and DES-EDE3-CBC for message encryption and SHA for message digest

RSA (in conjunction with the specified encryption and digest algorithm combinations) is used for both key generation and authentication on SSL connections. This usage is independent of whether or not a CA trustpoint is configured.

Secure Copy Protocol

The Secure Copy Protocol (SCP) feature provides a secure and authenticated method for copying switch configurations or switch image files. SCP relies on Secure Shell (SSH), an application and a protocol that provides a secure replacement for the Berkeley r-tools.

For SSH to work, the switch needs an RSA public/private key pair. This is the same with SCP, which relies on SSH for its secure transport.

Because SSH also relies on AAA authentication, and SCP relies further on AAA authorization, correct configuration is necessary.

- Before enabling SCP, you must correctly configure SSH, authentication, and authorization on the switch.
- Because SCP relies on SSH for its secure transport, the switch must have an Rivest, Shamir, and Adelman (RSA) key pair.

Note When using SCP, you cannot enter the password into the copy command. You must enter the password when prompted.

The behavior of SCP is similar to that of remote copy (rcp), which comes from the Berkeley r-tools suite, except that SCP relies on SSH for security. SCP also requires that authentication, authorization, and accounting (AAA) authorization be configured so the router can determine whether the user has the correct privilege level.

A user who has appropriate authorization can use SCP to copy any file in the Cisco IOS File System (IFS) to and from a switch by using the **copy** command. An authorized administrator can also do this from a workstation.

Note For information about how to configure and verify SCP, see the "Secure Copy Protocol" section in the *Cisco IOS Security Configuration Guide: Securing User Services, Release 12.4*: http://www.cisco.com/en/US/docs/ios/sec_user_services/configuration/guide/sec_secure_copy_ps6350_TSD_Products_Configuration_Guide_Chapter.html

How to Configure Switch-Based Authentication

Configuring Password Protection

Setting or Changing a Static Enable Password

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**enable password** *password*	Defines a new password or changes an existing password for access to privileged EXEC mode. By default, no password is defined. *password*—Specifies a string from 1 to 25 alphanumeric characters. The string cannot start with a number, is case sensitive, and allows spaces but ignores leading spaces. It can contain the question mark (?) character if you precede the question mark with the key combination Crtl-v when you create the password; for example, to create the password abc?123, do this: Enter **abc**. Press **Crtl-v**. Enter **?123**. When the system prompts you to enter the enable password, you need not precede the question mark by pressing Ctrl V; you can enter **abc?123** at the password prompt.
Step 3	**end**	Returns to privileged EXEC mode.

Protecting Enable and Enable Secret Passwords with Encryption

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**enable password** [**level** *level*] {*password* | *encryption-type encrypted-password*}	Defines a new password or changes an existing password for access to privileged EXEC mode.
	or	or
	enable secret [**level** *level*] {*password* | *encryption-type encrypted-password*}	Defines a secret password, which is saved using a nonreversible encryption method. • (Optional) *level*—Specifies the range is from 0 to 15. Level 1 is normal user EXEC mode privileges. The default level is 15 (privileged EXEC mode privileges). • *password*—Specifies a string from 1 to 25 alphanumeric characters. The string cannot start with a number, is case sensitive, and allows spaces but ignores leading spaces. By default, no password is defined. • (Optional) *encryption-type*—Only type 5, a Cisco proprietary encryption algorithm, is available. If you specify an encryption type, you must provide an encrypted password—an encrypted password that you copy from another switch configuration. **Note** If you specify an encryption type and then enter a clear text password, you cannot reenter privileged EXEC mode. You cannot recover a lost encrypted password by any method.
Step 3	**service password-encryption**	(Optional) Encrypts the password when the password is defined or when the configuration is written. Encryption prevents the password from being readable in the configuration file.
Step 4	**end**	Returns to privileged EXEC mode.

Disabling Password Recovery

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**no service password-recovery**	Disables password recovery. This setting is saved in an area of the flash memory that is accessible by the boot loader and the Cisco IOS image, but it is not part of the file system and is not accessible by any user.
Step 3	**end**	Returns to privileged EXEC mode.
Step 4	**show version**	Verifies the configuration by checking the last few lines of the command output.

Setting a Telnet Password for a Terminal Line

	Command	Purpose
Step 1		Attaches a PC or workstation with emulation software to the switch console port. The default data characteristics of the console port are 9600, 8, 1, no parity. You might need to press the Return key several times to see the command-line prompt.
Step 2	**enable password** *password*	Enters privileged EXEC mode.
Step 3	**configure terminal**	Enters global configuration mode.
Step 4	**line vty 0 15**	Configures the number of Telnet sessions (lines), and enters line configuration mode. There are 16 possible sessions on a command-capable switch. The 0 and 15 mean that you are configuring all 16 possible Telnet sessions.
Step 5	**password** *password*	Enters a Telnet password for the line or lines. *password*—Specifies a string from 1 to 25 alphanumeric characters. The string cannot start with a number, is case sensitive, and allows spaces but ignores leading spaces. By default, no password is defined.
Step 6	**end**	Returns to privileged EXEC mode.

Configuring Username and Password Pairs

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**username** *name* [**privilege** *level*] {**password** *encryption-type password*}	Enters the username, privilege level, and password for each user. • *name*—Specifies the user ID as one word. Spaces and quotation marks are not allowed. • (Optional) *level*—Specifies the privilege level the user has after gaining access. The range is 0 to 15. Level 15 gives privileged EXEC mode access. Level 1 gives user EXEC mode access. • *encryption-type*—Enters 0 to specify that an unencrypted password will follow. Enter 7 to specify that a hidden password will follow. • *password*—Specifies the password the user must enter to gain access to the switch. The password must be from 1 to 25 characters, can contain embedded spaces, and must be the last option specified in the **username** command. • To disable username authentication for a specific user, use the **no username** *name* global configuration command.
Step 3	**line console 0** or **line vty 0 15**	Enters line configuration mode, and configure the console port (line 0) or the VTY lines (line 0 to 15).

	Command	Purpose
Step 4	**login local**	Enables local password checking at login time. Authentication is based on the username specified in Step 2. To disable password checking and allow connections without a password, use the **no login** line configuration command.
Step 5	**end**	Returns to privileged EXEC mode.

Setting the Privilege Level for a Command

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**privilege** *mode* **level** *level command*	Sets the privilege level for a command. • *mode*—Enters **configure** for global configuration mode, **exec** for EXEC mode, **interface** for interface configuration mode, or **line** for line configuration mode. • *level*—The range is from 0 to 15. Level 1 is for normal user EXEC mode privileges. Level 15 is the level of access permitted by the **enable** password. • *command*—Specifies the command to which you want to restrict access.
Step 3	**enable password level** *level password*	Specifies the enable password for the privilege level. • *level*—The range is from 0 to 15. Level 1 is for normal user EXEC mode privileges. • *password*—Specifies a string from 1 to 25 alphanumeric characters. The string cannot start with a number, is case sensitive, and allows spaces but ignores leading spaces. By default, no password is defined.
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show privilege**	Verifies the password and accesses level configuration.

Changing the Default Privilege Level for Lines

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**line vty** *line*	Selects the virtual terminal line on which to restrict access.
Step 3	**privilege level** *level*	Changes the default privilege level for the line. *level*—The range is from 0 to 15. Level 1 is for normal user EXEC mode privileges. Level 15 is the level of access permitted by the **enable** password.
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show privilege**	Verifies the password and accesses level configuration.

Logging Into and Exiting a Privilege Level

Command	Purpose
enable *level*	Logs in to a specified privilege level. *level*—The range is 0 to 15.
disable *level*	Exits to a specified privilege level. *level*—The range is 0 to 15.

Configuring TACACS+

This section describes how to configure your switch to support TACACS+. At a minimum, you must identify the host or hosts maintaining the TACACS+ daemon and define the method lists for TACACS+ authentication. You can optionally define method lists for TACACS+ authorization and accounting. A method list defines the sequence and methods to be used to authenticate, to authorize, or to keep accounts on a user. You can use method lists to designate one or more security protocols to be used, thus ensuring a backup system if the initial method fails. The software uses the first method listed to authenticate, to authorize, or to keep accounts on users; if that method does not respond, the software selects the next method in the list. This process continues until there is successful communication with a listed method or the method list is exhausted.

Identifying the TACACS+ Server Host and Setting the Authentication Key

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**tacacs-server host** *hostname* [**port** *integer*] [**timeout** *integer*] [**key** *string*]	Identifies the IP host or hosts maintaining a TACACS+ server. Enters this command multiple times to create a list of preferred hosts. The software searches for hosts in the order in which you specify them. • *hostname*—Specifies the name or IP address of the host. • (Optional) **port** *integer*—Specifies a server port number. The default is port 49. The range is 1 to 65535. • (Optional) **timeout** *integer*—Specifies a time in seconds the switch waits for a response from the daemon before it times out and declares an error. The default is 5 seconds. The range is 1 to 1000 seconds. • (Optional) **key** *string*—Specifies the encryption key for encrypting and decrypting all traffic between the switch and the TACACS+ daemon. You must configure the same key on the TACACS+ daemon for encryption to be successful.
Step 3	**aaa new-model**	Enables AAA.
Step 4	**aaa group server tacacs+** *group-name*	(Optional) Defines the AAA server-group with a group name. This command puts the switch in a server group subconfiguration mode.

	Command	Purpose
Step 5	**server** *ip-address*	(Optional) Associates a particular TACACS+ server with the defined server group. Repeat this step for each TACACS+ server in the AAA server group. Each server in the group must be previously defined in Step 2.
Step 6	**end**	Returns to privileged EXEC mode.
Step 7	**show tacacs**	Verifies your entries.

Configuring TACACS+ Login Authentication

Before You Begin

To secure the switch for HTTP access by using AAA methods, you must configure the switch with the **ip http authentication aaa** global configuration command. Configuring AAA authentication does not secure the switch for HTTP access by using AAA methods.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**aaa new-model**	Enables AAA.

	Command	Purpose
Step 3	**aaa authentication login** {**default** | *list-name*} *method1* [*method2...*]	Creates a login authentication method list. • To create a default list that is used when a named list is *not* specified in the **login authentication** command, use the **default** keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all ports. • *list-name*—Specifies a character string to name the list you are creating. • *method1...*—Specifies the actual method the authentication algorithm tries. The additional methods of authentication are used only if the previous method returns an error, not if it fails. Select one of these methods: • **enable**—Uses the enable password for authentication. Before you can use this authentication method, you must define an enable password by using the **enable** *password* global configuration command. • **group tacacs+**—Uses TACACS+ authentication. Before you can use this authentication method, you must configure the TACACS+ server. For more information, see the "Identifying the TACACS+ Server Host and Setting the Authentication Key" section on page 12-30. • **line**—Uses the line password for authentication. Before you can use this authentication method, you must define a line password. Use the **password** *password* line configuration command. • **local**—Uses the local username database for authentication. You must enter username information in the database. Use the **username** *password* global configuration command. • **local-case**—Uses a case-sensitive local username database for authentication. You must enter username information in the database by using the **username** *name* **password** global configuration command. • **none**—Does not use any authentication for login.
Step 4	**line** [**console** | **tty** | **vty**] *line-number* [*ending-line-number*]	Enters line configuration mode, and configures the lines to which you want to apply the authentication list.
Step 5	**login authentication** {**default** | *list-name*}	Applies the authentication list to a line or set of lines. • If you specify **default**, use the default list created with the **aaa authentication login** command. • *list-name*—Specifies the list created with the **aaa authentication login** command.
Step 6	**end**	Returns to privileged EXEC mode.

Configuring TACACS+ Authorization for Privileged EXEC Access and Network Services

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**aaa authorization network tacacs+**	Configures the switch for user TACACS+ authorization for all network-related service requests.
Step 3	**aaa authorization exec tacacs+**	Configures the switch for user TACACS+ authorization if the user has privileged EXEC access. The **exec** keyword might return user profile information (such as **autocommand** information).
Step 4	**end**	Returns to privileged EXEC mode.

Starting TACACS+ Accounting

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**aaa accounting network start-stop tacacs+**	Enables TACACS+ accounting for all network-related service requests.
Step 3	**aaa accounting exec start-stop tacacs+**	Enables TACACS+ accounting to send a start-record accounting notice at the beginning of a privileged EXEC process and a stop-record at the end.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring Radius Server Communication

Before You Begin

You should have access to and should configure a RADIUS server before configuring RADIUS features on your switch.

At a minimum, you must identify the host or hosts that run the RADIUS server software and define the method lists for RADIUS authentication. You can optionally define method lists for RADIUS authorization and accounting.

Some configuration settings need to be configured on the RADIUS server that include the IP address of the switch and the key string to be shared by both the server and the switch.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**radius-server host** {*hostname* | *ip-address*} [**auth-port** *port-number*] [**acct-port** *port-number*] [**timeout** *seconds*] [**retransmit** *retries*] [**key** *string*]	Specifies the IP address or hostname of the remote RADIUS server host. • (Optional) **auth-port** *port-number*—Specifies the UDP destination port for authentication requests. • (Optional) **acct-port** *port-number*—Specifies the UDP destination port for accounting requests. • (Optional) **timeout** *seconds*—Specifies the time interval that the switch waits for the RADIUS server to reply before resending. The range is 1 to 1000. This setting overrides the **radius-server timeout** global configuration command setting. If no timeout is set with the **radius-server host** command, the setting of the **radius-server timeout** command is used. • (Optional) **retransmit** *retries*—Specifies the number of times a RADIUS request is resent to a server if that server is not responding or responding slowly. The range is 1 to 1000. If no retransmit value is set with the **radius-server host** command, the setting of the **radius-server retransmit** global configuration command is used. • (Optional) **key** *string*—Specifies the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. **Note** The key is a text string that must match the encryption key used on the RADIUS server. Always configure the key as the last item in the **radius-server host** command. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key. To configure the switch to recognize more than one host entry associated with a single IP address, enter this command as many times as necessary, making sure that each UDP port number is different. The switch software searches for hosts in the order in which you specify them. Set the timeout, retransmit, and encryption key values to use with the specific RADIUS host.
Step 3	**end**	Returns to privileged EXEC mode.

Defining AAA Server Groups

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**radius-server host** {*hostname* | *ip-address*} [**auth-port** *port-number*] [**acct-port** *port-number*] [**timeout** *seconds*] [**retransmit** *retries*] [**key** *string*]	Specifies the IP address or hostname of the remote RADIUS server host. • (Optional) **auth-port** *port-number*—Specifies the UDP destination port for authentication requests. • (Optional) **acct-port** *port-number*—Specifies the UDP destination port for accounting requests. • (Optional) **timeout** *seconds*—Specifies the time interval that the switch waits for the RADIUS server to reply before resending. The range is 1 to 1000. This setting overrides the **radius-server timeout** global configuration command setting. If no timeout is set with the **radius-server host** command, the setting of the **radius-server timeout** command is used. • (Optional) **retransmit** *retries*—Specifies the number of times a RADIUS request is resent to a server if that server is not responding or responding slowly. The range is 1 to 1000. If no retransmit value is set with the **radius-server host** command, the setting of the **radius-server retransmit** global configuration command is used. • (Optional) **key** *string*, specifies the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. **Note** The key is a text string that must match the encryption key used on the RADIUS server. Always configure the key as the last item in the **radius-server host** command. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key. To configure the switch to recognize more than one host entry associated with a single IP address, enter this command as many times as necessary, making sure that each UDP port number is different. The switch software searches for hosts in the order in which you specify them. Set the timeout, retransmit, and encryption key values to use with the specific RADIUS host.
Step 3	**aaa new-model**	Enables AAA.
Step 4	**aaa group server radius** *group-name*	Defines the AAA server group with a group name. This command puts the switch in a server group configuration mode.
Step 5	**server** *ip-address*	Associates a particular RADIUS server with the defined server group. Repeat this step for each RADIUS server in the AAA server group. Each server in the group must be previously defined in Step 2.
Step 6	**end**	Returns to privileged EXEC mode.
Step 7		Enable RADIUS login authentication. See the "Defining AAA Server Groups" section on page 12-35.

Configuring RADIUS Login Authentication

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**aaa new-model**	Enables AAA.
Step 3	**aaa authentication login** {**default** | *list-name*} *method1* [*method2...*]	Creates a login authentication method list. • To create a default list that is used when a named list is *not* specified in the **login authentication** command, use the **default** keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all ports. • *list-name*—Specifies a character string to name the list you are creating. • *method1...*—Specifies the actual method the authentication algorithm tries. The additional methods of authentication are used only if the previous method returns an error, not if it fails. Select one of these methods: – **enable**—Uses the enable password for authentication. Before you can use this authentication method, you must define an enable password by using the **enable** *password* global configuration command. – **group radius**—Uses RADIUS authentication. Before you can use this authentication method, you must configure the RADIUS server. For more information, see the "RADIUS Server Host" section on page 12-14. – **line**—Uses the line password for authentication. Before you can use this authentication method, you must define a line password. Use the **password** *password* line configuration command. – **local**—Uses the local username database for authentication. You must enter username information in the database. Use the **username** *name* **password** global configuration command. – **local-case**—Uses a case-sensitive local username database for authentication. You must enter username information in the database by using the **username** *password* global configuration command. – **none**—Does not use any authentication for login.
Step 4	**line** [**console** | **tty** | **vty**] *line-number* [*ending-line-number*]	Enters line configuration mode, and configures the lines to which you want to apply the authentication list.
Step 5	**login authentication** {**default** | *list-name*}	Applies the authentication list to a line or set of lines. • If you specify **default**, use the default list created with the **aaa authentication login** command. • *list-name*—Specifies the list created with the **aaa authentication login** command.
Step 6	**end**	Returns to privileged EXEC mode.

Configuring RADIUS Authorization for User Privileged Access and Network Services

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**aaa authorization network radius**	Configures the switch for user RADIUS authorization for all network-related service requests.
Step 3	**aaa authorization exec radius**	Configures the switch for user RADIUS authorization if the user has privileged EXEC access. The **exec** keyword might return user profile information (such as **autocommand** information).
Step 4	**end**	Returns to privileged EXEC mode.

Starting RADIUS Accounting

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**aaa accounting network start-stop radius**	Enables RADIUS accounting for all network-related service requests.
Step 3	**aaa accounting exec start-stop radius**	Enables RADIUS accounting to send a start-record accounting notice at the beginning of a privileged EXEC process and a stop-record at the end.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring Settings for All RADIUS Servers

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**radius-server key** *string*	Specifies the shared secret text string used between the switch and all RADIUS servers. **Note** The key is a text string that must match the encryption key used on the RADIUS server. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.
Step 3	**radius-server retransmit** *retries*	Specifies the number of times the switch sends each RADIUS request to the server before giving up. The default is 3; the range 1 to 1000.
Step 4	**radius-server timeout** *seconds*	Specifies the number of seconds a switch waits for a reply to a RADIUS request before resending the request. The default is 5 seconds; the range is 1 to 1000.
Step 5	**radius-server deadtime** *minutes*	Specifies the number of minutes a RADIUS server, which is not responding to authentication requests, to be skipped, thus avoiding the wait for the request to timeout before trying the next configured server. The default is 0; the range is 1 to 1440 minutes.

	Command	Purpose
Step 6	**radius-server vsa send** [**accounting** | **authentication**]	Enables the switch to recognize and use VSAs as defined by RADIUS IETF attribute 26. • (Optional) **accounting**—Limits the set of recognized vendor-specific attributes to only accounting attributes. • (Optional) **authentication**—Limits the set of recognized vendor-specific attributes to only authentication attributes. If you enter this command without keywords, both accounting and authentication vendor-specific attributes are used.
Step 7	**end**	Returns to privileged EXEC mode.

Configuring the Switch for Vendor-Proprietary RADIUS Server Communication

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**radius-server host** {*hostname* | *ip-address*} **non-standard**	Specifies the IP address or hostname of the remote RADIUS server host and identifies that it is using a vendor-proprietary implementation of RADIUS.
Step 3	**radius-server key** *string*	Specifies the shared secret text string used between the switch and the vendor-proprietary RADIUS server. The switch and the RADIUS server use this text string to encrypt passwords and exchange responses. **Note** The key is a text string that must match the encryption key used on the RADIUS server. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show running-config**	Verifies your settings.
Step 6	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring CoA on the Switch

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**aaa new-model**	Enables AAA.
Step 3	**aaa server radius dynamic-author**	Configures the switch as an authentication, authorization, and accounting (AAA) server to facilitate interaction with an external policy server.

	Command	Purpose
Step 4	**client** {*ip-address* | *name*} [**vrf** *vrfname*] [**server-key** *string*]	Enters dynamic authorization local server configuration mode and specifies a RADIUS client from which a device will accept CoA and disconnect requests.
Step 5	**server-key** [**0** | **7**] *string*	Configures the RADIUS key to be shared between a device and RADIUS clients.
Step 6	**port** *port-number*	Specifies the port on which a device listens for RADIUS requests from configured RADIUS clients.
Step 7	**auth-type** {**any** | **all** | **session-key**}	Specifies the type of authorization the switch uses for RADIUS clients. The client must match all the configured attributes for authorization.
Step 8	**ignore session-key**	(Optional) Configures the switch to ignore the session-key.
Step 9	**ignore server-key**	(Optional) Configures the switch to ignore the server-key.
Step 10	**authentication command bounce-port ignore**	(Optional) Configures the switch to ignore a CoA request to temporarily disable the port hosting a session. The purpose of temporarily disabling the port is to trigger a DHCP renegotiation from the host when a VLAN change occurs and there is no supplicant on the endpoint to detect the change.
Step 11	**authentication command disable-port ignore**	(Optional) Configures the switch to ignore a nonstandard command requesting that the port hosting a session be administratively shut down. Shutting down the port results in termination of the session. Uses standard CLI or SNMP commands to reenable the port.
Step 12	**end**	Returns to privileged EXEC mode.

Configuring the Switch for Local Authentication and Authorization

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**aaa new-model**	Enables AAA.
Step 3	**aaa authentication login default local**	Sets the login authentication to use the local username database. The **default** keyword applies the local user database authentication to all ports.
Step 4	**aaa authorization exec local**	Configures user AAA authorization, checks the local database, and allows the user to run an EXEC shell.
Step 5	**aaa authorization network local**	Configures user AAA authorization for all network-related service requests.

	Command	Purpose
Step 6	**username** *name* [**privilege** *level*] {**password** *encryption-type password*}	Enters the local database, and establishes a username-based authentication system. Repeat this command for each user. • *name*—Specifies the user ID as one word. Spaces and quotation marks are not allowed. • (Optional) *level*—Specifies the privilege level the user has after gaining access. The range is 0 to 15. Level 15 gives privileged EXEC mode access. Level 0 gives user EXEC mode access. • *encryption-type*—Enters 0 to specify that an unencrypted password follows. Enter 7 to specify that a hidden password follows. • *password*—Specifies the password the user must enter to gain access to the switch. The password must be from 1 to 25 characters, can contain embedded spaces, and must be the last option specified in the **username** command.
Step 7	**end**	Returns to privileged EXEC mode.
Step 8	**show running-config**	Verifies your entries.
Step 9	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring Secure Shell

Setting Up the Switch to Run SSH

	Task	Purpose
Step 1	Download the cryptographic software image from Cisco.com.	(Required) For more information, see the notes for this release.
Step 2	Configure a hostname and IP domain name for the switch.	Follow this procedure only if you are configuring the switch as an SSH server.
Step 3	Generate an RSA key pair for the switch, which automatically enables SSH.	Follow this procedure only if you are configuring the switch as an SSH server.
Step 4	Configure user authentication for local or remote access.	(Required) For more information, see the "Configuring the Switch for Local Authentication and Authorization" section on page 12-39.

Configuring the SSH Server

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**hostname** *hostname*	Configures a hostname for your switch.
Step 3	**ip domain-name** *domain_name*	Configures a host domain for your switch.

	Command	Purpose
Step 4	**crypto key generate rsa**	Enables the SSH server for local and remote authentication on the switch and generates an RSA key pair. We recommend that a minimum modulus size of 1024 bits. When you generate RSA keys, you are prompted to enter a modulus length. A longer modulus length might be more secure, but it takes longer to generate and to use.
Step 5	**ip ssh version** [**1** | **2**]	(Optional) Configures the switch to run SSH Version 1 or SSH Version 2. • **1**—Configures the switch to run SSH Version 1. • **2**—Configures the switch to run SSH Version 2. If you do not enter this command or do not specify a keyword, the SSH server selects the latest SSH version supported by the SSH client. For example, if the SSH client supports SSHv1 and SSHv2, the SSH server selects SSHv2.
Step 6	**ip ssh** {**timeout** *seconds* | **authentication-retries** *number*}	Configures the SSH control parameters. • Specifies the time-out value in seconds; the default is 120 seconds. The range is 0 to 120 seconds. This parameter applies to the SSH negotiation phase. After the connection is established, the switch uses the default time-out values of the CLI-based sessions. By default, up to five simultaneous, encrypted SSH connections for multiple CLI-based sessions over the network are available (session 0 to session 4). After the execution shell starts, the CLI-based session time-out value returns to the default of 10 minutes. • Specifies the number of times that a client can reauthenticate to the server. The default is 3; the range is 0 to 5. Repeat this step when configuring both parameters.
Step 7	**line vty** *line_number* [*ending_line_number*] **transport input ssh**	(Optional) Configures the virtual terminal line settings. • Enters line configuration mode to configure the virtual terminal line settings. *line_number* and *ending_line_number* specifiy a pair of lines. The range is 0 to 15. • Specifies that the switch prevent non-SSH Telnet connections. This limits the router to only SSH connections.
Step 8	**end**	Returns to privileged EXEC mode.
Step 9	**show ip ssh** or **show ssh**	Shows the version and configuration information for your SSH server. Shows the status of the SSH server on the switch.

Configuring Secure HTTP Servers and Clients

Configuring a CA Trustpoint

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**hostname** *hostname*	Specifies the hostname of the switch (required only if you have not previously configured a hostname).
Step 3	**ip domain-name** *domain-name*	Specifies the IP domain name of the switch (required only if you have not previously configured an IP domain name).
Step 4	**crypto key generate rsa**	(Optional) Generates an RSA key pair. RSA key pairs are required before you can obtain a certificate for the switch. RSA key pairs are generated automatically. You can use this command to regenerate the keys, if needed.
Step 5	**crypto ca trustpoint** *name*	Specifies a local configuration name for the CA trustpoint and enter CA trustpoint configuration mode.
Step 6	**enrollment url** *url*	Specifies the URL to which the switch should send certificate requests.
Step 7	**enrollment http-proxy** *host-name port-number*	(Optional) Configures the switch to obtain certificates from the CA through an HTTP proxy server.
Step 8	**crl query** *url*	Configures the switch to request a certificate revocation list (CRL) to ensure that the certificate of the peer has not been revoked.
Step 9	**primary**	(Optional) Specifies that the trustpoint should be used as the primary (default) trustpoint for CA requests.
Step 10	**exit**	Exits CA trustpoint configuration mode and returns to global configuration mode.
Step 11	**crypto ca authentication** *name*	Authenticates the CA by getting the public key of the CA. Uses the same name used in Step 5.
Step 12	**crypto ca enroll** *name*	Obtains the certificate from the specified CA trustpoint. This command requests a signed certificate for each RSA key pair.
Step 13	**end**	Returns to privileged EXEC mode.
Step 14	**show crypto ca trustpoints**	Verifies the configuration.

Configuring the Secure HTTP Server

Before You Begin

If you are using a certificate authority for certification, you should use the previous procedure to configure the CA trustpoint on the switch before enabling the HTTP server. If you have not configured a CA trustpoint, a self-signed certificate is generated the first time that you enable the secure HTTP server. After you have configured the server, you can configure options (path, access list to apply, maximum number of connections, or timeout policy) that apply to both standard and secure HTTP servers.

	Command	Purpose

Step 1	**show ip http server status**	(Optional) Displays the status of the HTTP server to determine if the secure HTTP server feature is supported in the software. You should see one of these lines in the output:
		`HTTP secure server capability: Present`
		or
		`HTTP secure server capability: Not present`

Step 2	**configure terminal**	Enters global configuration mode.
Step 3	**ip http secure-server**	Enables the HTTPS server if it has been disabled. The HTTPS server is enabled by default.
Step 4	**ip http secure-port** *port-number*	(Optional) Specifies the port number to be used for the HTTPS server. The default port number is 443. Valid options are 443 or any number in the range 1025 to 65535.
Step 5	**ip http secure-ciphersuite** {[**3des-ede-cbc-sha**] [**rc4-128-md5**] [**rc4-128-sha**] [**des-cbc-sha**]}	(Optional) Specifies the CipherSuites (encryption algorithms) to be used for encryption over the HTTPS connection. If you do not have a reason to specify a particularly CipherSuite, you should allow the server and client to negotiate a CipherSuite that they both support. This is the default.
Step 6	**ip http secure-client-auth**	(Optional) Configures the HTTP server to request an X.509v3 certificate from the client for authentication during the connection process. The default is for the client to request a certificate from the server, but the server does not attempt to authenticate the client.
Step 7	**ip http secure-trustpoint** *name*	Specifies the CA trustpoint to use to get an X.509v3 security certificate and to authenticate the client certificate connection. **Note** Use of this command assumes you have already configured a CA trustpoint according to the previous procedure.
Step 8	**ip http path** *path-name*	(Optional) Sets a base HTTP path for HTML files. The path specifies the location of the HTTP server files on the local system (usually located in system flash memory).
Step 9	**ip http access-class** *access-list-number*	(Optional) Specifies an access list to use to allow access to the HTTP server.
Step 10	**ip http max-connections** *value*	(Optional) Sets the maximum number of concurrent connections that are allowed to the HTTP server. The range is 1 to 16; the default value is 5.
Step 11	**ip http timeout-policy idle** *seconds* **life** *seconds* **requests** *value*	(Optional) Specifies how long a connection to the HTTP server can remain open under the defined circumstances: • **idle**—Specifies the maximum time period when no data is received or response data cannot be sent. The range is 1 to 600 seconds. The default is 180 seconds (3 minutes). • **life**—Specifies the maximum time period from the time that the connection is established. The range is 1 to 86400 seconds (24 hours). The default is 180 seconds. • **requests**—Specifies the maximum number of requests processed on a persistent connection. The maximum value is 86400. The default is 1.

	Command	Purpose
Step 12	**end**	Returns to privileged EXEC mode.
Step 13	**show ip http server secure status**	Displays the status of the HTTP secure server to verify the configuration.

Configuring the Secure HTTP Client

Before You Begin

The standard HTTP client and secure HTTP client are always enabled. A certificate authority is required for secure HTTP client certification. This procedure assumes that you have previously configured a CA trustpoint on the switch. If a CA trustpoint is not configured and the remote HTTPS server requires client authentication, connections to the secure HTTP client fail.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip http client secure-trustpoint** *name*	(Optional) Specifies the CA trustpoint to be used if the remote HTTP server requests client authentication. Using this command assumes that you have already configured a CA trustpoint by using the previous procedure. The command is optional if client authentication is not needed or if a primary trustpoint has been configured.
Step 3	**ip http client secure-ciphersuite** {[**3des-ede-cbc-sha**] [**rc4-128-md5**] [**rc4-128-sha**] [**des-cbc-sha**]}	(Optional) Specifies the CipherSuites (encryption algorithms) to be used for encryption over the HTTPS connection. If you do not have a reason to specify a particular CipherSuite, you should allow the server and client to negotiate a CipherSuite that they both support. This is the default.
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show ip http client secure status**	Displays the status of the HTTP secure server to verify the configuration.
Step 6	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Monitoring and Maintaining Switch-Based Authentication

Command	Purpose
show running-config	Verifies your configured entries.
copy running-config startup-config	Saves your entries in the configuration file.
show tacacs	Displays the TACACS+ server statistics.
debug radius	Displays the information associated with RADIUS.
debug aaa coa	Displays the debug information for CoA processing.
debug cmdhd	Displays the debug information for the command handler.
show aaa attributes protocol radius	Displays the RADIUS attributes.
show ip ssh	Displays the version and configuration information for the SSH server.
show ssh	Displays the status of the SSH server.

Command	Purpose
show ip http client secure status	Displays the HTTP secure client configuration.
show ip http server secure status	Displays the HTTP secure server configuration.

Configuration Examples for Configuring Switch-Based Authentication

Changing the Enable Password: Example

This example shows how to change the enable password to *l1u2c3k4y5*. The password is not encrypted and provides access to level 15 (traditional privileged EXEC mode access):

```
Switch(config)# enable password l1u2c3k4y5
```

Configuring the Encrypted Password: Example

This example shows how to configure the encrypted password *1FaD0$Xyti5Rkls3LoyxzS8* for privilege level 2:

```
Switch(config)# enable secret level 2 5 $1$FaD0$Xyti5Rkls3LoyxzS8
```

Setting the Telnet Password for a Terminal Line: Example

This example shows how to set the Telnet password to *let45me67in89*:

```
Switch(config)# line vty 10
Switch(config-line)# password let45me67in89
```

Setting the Privilege Level for a Command: Example

This example shows how to set the **configure** command to privilege level 14 and define *SecretPswd14* as the password users must enter to use level 14 commands:

```
Switch(config)# privilege exec level 14 configure
Switch(config)# enable password level 14 SecretPswd14
```

Configuring the RADIUS Server: Examples

This example shows how to configure one RADIUS server to be used for authentication and another to be used for accounting:

```
Switch(config)# radius-server host 172.29.36.49 auth-port 1612 key rad1
Switch(config)# radius-server host 172.20.36.50 acct-port 1618 key rad2
```

This example shows how to configure *host1* as the RADIUS server and to use the default ports for both authentication and accounting:

```
Switch(config)# radius-server host host1
```

Defining AAA Server Groups: Example

In this example, the switch is configured to recognize two different RADIUS group servers (*group1* and *group2*). Group1 has two different host entries on the same RADIUS server configured for the same services. The second host entry acts as a fail-over backup to the first entry.

```
Switch(config)# radius-server host 172.20.0.1 auth-port 1000 acct-port 1001
Switch(config)# radius-server host 172.10.0.1 auth-port 1645 acct-port 1646
Switch(config)# aaa new-model
Switch(config)# aaa group server radius group1
Switch(config-sg-radius)# server 172.20.0.1 auth-port 1000 acct-port 1001
Switch(config-sg-radius)# exit
Switch(config)# aaa group server radius group2
Switch(config-sg-radius)# server 172.20.0.1 auth-port 2000 acct-port 2001
Switch(config-sg-radius)# exit
```

Configuring Vendor-Specific RADIUS Attributes: Examples

This example shows how to provide a user logging in from a switch with immediate access to privileged EXEC commands:

```
cisco-avpair= "shell:priv-lvl=15"
```

This example shows how to specify an authorized VLAN in the RADIUS server database:

```
cisco-avpair= "tunnel-type(#64)=VLAN(13)"
cisco-avpair= "tunnel-medium-type(#65)=802 media(6)"
cisco-avpair= "tunnel-private-group-id(#81)=vlanid"
```

This example shows how to apply an input ACL in ASCII format to an interface for the duration of this connection:

```
cisco-avpair= "ip:inacl#1=deny ip 10.10.10.10 0.0.255.255 20.20.20.20 255.255.0.0"
cisco-avpair= "ip:inacl#2=deny ip 10.10.10.10 0.0.255.255 any"
cisco-avpair= "mac:inacl#3=deny any any decnet-iv"
```

This example shows how to apply an output ACL in ASCII format to an interface for the duration of this connection:

```
cisco-avpair= "ip:outacl#2=deny ip 10.10.10.10 0.0.255.255 any"
```

Configuring a Vendor-Proprietary RADIUS Host: Example

This example shows how to specify a vendor-proprietary RADIUS host and to use a secret key of *rad124* between the switch and the server:

```
Switch(config)# radius-server host 172.20.30.15 nonstandard
Switch(config)# radius-server key rad124
```

Sample Output for a Self-Signed Certificate: Example

If a self-signed certificate has been generated, this information is included in the output of the **show running-config** privileged EXEC command. This is a partial sample output from that command displaying a self-signed certificate.

```
Switch# show running-config
Building configuration...
```

```
<output truncated>

crypto pki trustpoint TP-self-signed-3080755072
 enrollment selfsigned
 subject-name cn=IOS-Self-Signed-Certificate-3080755072
 revocation-check none
 rsakeypair TP-self-signed-3080755072
!
!
crypto ca certificate chain TP-self-signed-3080755072
 certificate self-signed 01
  3082029F 30820208 A0030201 02020101 300D0609 2A864886 F70D0101 04050030
  59312F30 2D060355 04031326 494F532D 53656C66 2D536967 6E65642D 43657274
  69666963 6174652D 33303830 37353530 37323126 30240609 2A864886 F70D0109

<output truncated>
```

You can remove this self-signed certificate by disabling the secure HTTP server and entering the **no crypto pki trustpoint TP-self-signed-30890755072** global configuration command. If you later reenable a secure HTTP server, a new self-signed certificate is generated.

Verifying Secure HTTP Connection: Example

To verify the secure HTTP connection by using a Web browser, enter https://*URL*, where the *URL* is the IP address or hostname of the server switch. If you configure a port other than the default port, you must also specify the port number after the URL. For example:

https://209.165.129:1026
or

https://host.domain.com:1026

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
Secure Copy Protocol configuration	*Cisco IOS Security Configuration Guide: Securing User Services*
RADIUS Server Load Balancing configuration	*Cisco IOS Security Configuration Guide*
Kerberos configuration examples	*Cisco IOS Security Configuration Guide: Security Server Protocols*
Authenticating a network service	*Cisco IOS Security Configuration Guide: Security Server Protocols*
Authenticating for KDC	*Cisco IOS Security Configuration Guide: Security Server Protocols*
Kerberos configuration task list	*Cisco IOS Security Configuration Guide: Security Server Protocols*
Login enhancement configuration	*Cisco IOS User Security Configuration Guide*

Related Topic	Document Title
Password protection commands	*Cisco IOS Security Command Reference*
Kerberos commands	*Cisco IOS Security Command Reference*
Secure Shell commands	*Cisco IOS Security Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 13

Configuring IEEE 802.1x Port-Based Authentication

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Restrictions for Configuring IEEE 802.1x Port-Based Authentication

- To use this feature, the switch must be running the LAN Base image.

Information About Configuring IEEE 802.1x Port-Based Authentication

IEEE 802.1x Port-Based Authentication

The standard defines a client-server-based access control and authentication protocol to prevent unauthorized clients from connecting to a LAN through publicly accessible ports. The authentication server authenticates each client connected to a switch port before making available any switch or LAN services.

Until the client is authenticated, IEEE 802.1x access control allows only Extensible Authentication Protocol over LAN (EAPOL), Cisco Discovery Protocol (CDP), and Spanning Tree Protocol (STP) traffic through the port to which the client is connected. After authentication, normal traffic passes through the port.

Device Roles

Figure 13-1 802.1x Device Roles

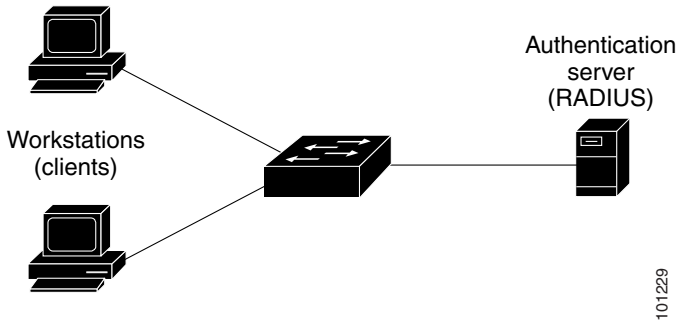

- *Client*—The device (workstation) that requests access to the LAN and switch services and responds to requests from the switch. The workstation must be running 802.1x-compliant client software such as that offered in the Microsoft Windows XP operating system. (The client is the *supplicant* in the 802.1x standard.)

  Note To resolve Windows XP network connectivity and 802.1x authentication issues, read the Microsoft Knowledge Base article: http://support.microsoft.com/support/kb/articles/Q303/5/97.ASP

- *Authentication server*—Performs the actual authentication of the client. The authentication server validates the identity of the client and notifies the switch whether or not the client is authorized to access the LAN and switch services. Because the switch acts as the proxy, the authentication service is transparent to the client. In this release, the RADIUS security system with Extensible Authentication Protocol (EAP) extensions is the only supported authentication server. It is available in Cisco Secure Access Control Server Version 3.0 or later. RADIUS operates in a client/server model in which secure authentication information is exchanged between the RADIUS server and one or more RADIUS clients.

- *Switch* (edge switch or wireless access point)—Controls the physical access to the network based on the authentication status of the client. The switch acts as an intermediary (proxy) between the client and the authentication server, requesting identity information from the client, verifying that information with the authentication server, and relaying a response to the client. The switch includes the RADIUS client, which is responsible for encapsulating and decapsulating the EAP frames and interacting with the authentication server. (The switch is the *authenticator* in the 802.1x standard.)

  When the switch receives EAPOL frames and relays them to the authentication server, the Ethernet header is stripped, and the remaining EAP frame is reencapsulated in the RADIUS format. The EAP frames are not modified during encapsulation, and the authentication server must support EAP within the native frame format. When the switch receives frames from the authentication server, the server's frame header is removed, leaving the EAP frame, which is then encapsulated for Ethernet and sent to the client.

  The devices that can act as intermediaries include the Cisco IE 2000, the Catalyst 3750-E, Catalyst 3560-E, Catalyst 3750, Catalyst 3560, Catalyst 3550, Catalyst 2975, Catalyst 2970, Catalyst 2960, Catalyst 2955, Catalyst 2950, Catalyst 2940 switches, or a wireless access point. These devices must be running software that supports the RADIUS client and 802.1x authentication.

Authentication Process

When 802.1x port-based authentication is enabled and the client supports 802.1x-compliant client software, these events occur:

- If the client identity is valid and the 802.1x authentication succeeds, the switch grants the client access to the network.
- If 802.1x authentication times out while waiting for an EAPOL message exchange and MAC authentication bypass is enabled, the switch can use the client MAC address for authorization. If the client MAC address is valid and the authorization succeeds, the switch grants the client access to the network. If the client MAC address is invalid and the authorization fails, the switch assigns the client to a guest VLAN that provides limited services if a guest VLAN is configured.
- If the switch gets an invalid identity from an 802.1x-capable client and a restricted VLAN is specified, the switch can assign the client to a restricted VLAN that provides limited services.
- If the RADIUS authentication server is unavailable (down) and inaccessible authentication bypass is enabled, the switch grants the client access to the network by putting the port in the critical-authentication state in the RADIUS-configured or the user-specified access VLAN.

Note Inaccessible authentication bypass is also referred to as critical authentication or the AAA fail policy.

Figure 13-2 Authentication Flowchart

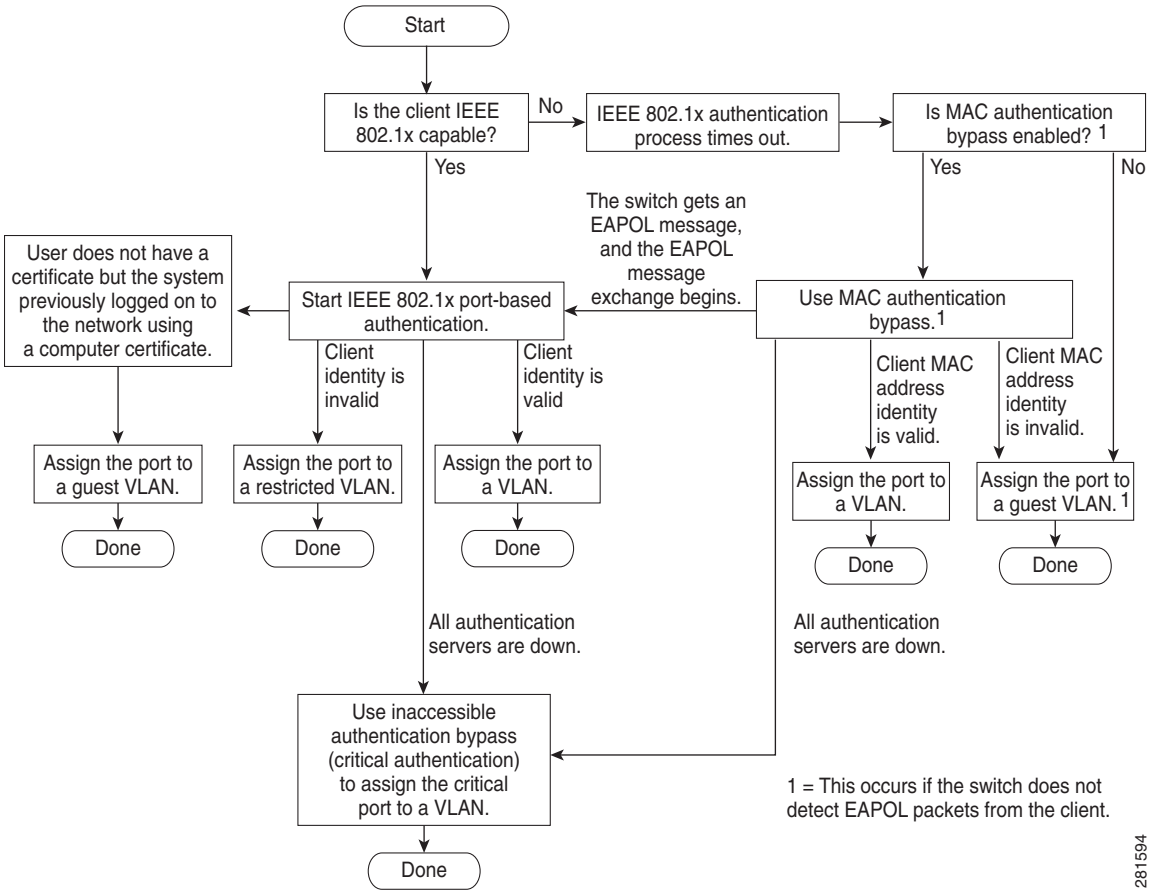

The switch reauthenticates a client when one of these situations occurs:

- Periodic reauthentication is enabled, and the reauthentication timer expires.

  You can configure the reauthentication timer to use a switch-specific value or to be based on values from the RADIUS server.

  After 802.1x authentication using a RADIUS server is configured, the switch uses timers based on the Session-Timeout RADIUS attribute (Attribute[27]) and the Termination-Action RADIUS attribute (Attribute [29]).

  The Session-Timeout RADIUS attribute (Attribute[27]) specifies the time after which reauthentication occurs.

  The Termination-Action RADIUS attribute (Attribute [29]) specifies the action to take during reauthentication. The actions are *Initialize* and *ReAuthenticate*. When the *Initialize* action is set (the attribute value is *DEFAULT*), the 802.1x session ends, and connectivity is lost during reauthentication. When the *ReAuthenticate* action is set (the attribute value is RADIUS-Request), the session is not affected during reauthentication.

- You manually reauthenticate the client by entering the **dot1x re-authenticate interface** *interface-id* privileged EXEC command.

If multidomain authentication (MDA) is enabled on a port, this flow can be used with some exceptions that are applicable to voice authorization. For more information on MDA, see the "Multidomain Authentication" section on page 13-10.

Switch-to-RADIUS-Server Communication

RADIUS security servers are identified by their hostname or IP address, hostname and specific UDP port numbers, or IP address and specific UDP port numbers. The combination of the IP address and the UDP port number creates a unique identifier, which enables RADIUS requests to be sent to multiple UDP ports on a server at the same IP address. If two different host entries on the same RADIUS server are configured for the same service—for example, authentication—the second host entry configured acts as the failover backup to the first one. The RADIUS host entries are tried in the order in which they were configured.

Authentication Initiation and Message Exchange

During 802.1x authentication, the switch or the client can initiate authentication. If you enable authentication on a port by using the **authentication port-control auto** interface configuration command, the switch initiates authentication when the link state changes from down to up or periodically as long as the port remains up and unauthenticated. The switch sends an EAP-request/identity frame to the client to request its identity. Upon receipt of the frame, the client responds with an EAP-response/identity frame.

However, if during boot up, the client does not receive an EAP-request/identity frame from the switch, the client can initiate authentication by sending an EAPOL-start frame, which prompts the switch to request the client's identity.

Note If 802.1x authentication is not enabled or supported on the network access device, any EAPOL frames from the client are dropped. If the client does not receive an EAP-request/identity frame after three attempts to start authentication, the client sends frames as if the port is in the authorized state. A port in the authorized state effectively means that the client has been successfully authenticated. For more information, see the "Ports in Authorized and Unauthorized States" section on page 13-9.

When the client supplies its identity, the switch begins its role as the intermediary, passing EAP frames between the client and the authentication server until authentication succeeds or fails. If the authentication succeeds, the switch port becomes authorized. If the authentication fails, authentication can be retried, the port might be assigned to a VLAN that provides limited services, or network access is not granted. For more information, see the "Ports in Authorized and Unauthorized States" section on page 13-9.

The specific exchange of EAP frames depends on the authentication method being used. Figure 13-3 shows a message exchange initiated by the client when the client uses the One-Time-Password (OTP) authentication method with a RADIUS server.

Figure 13-3 Message Exchange

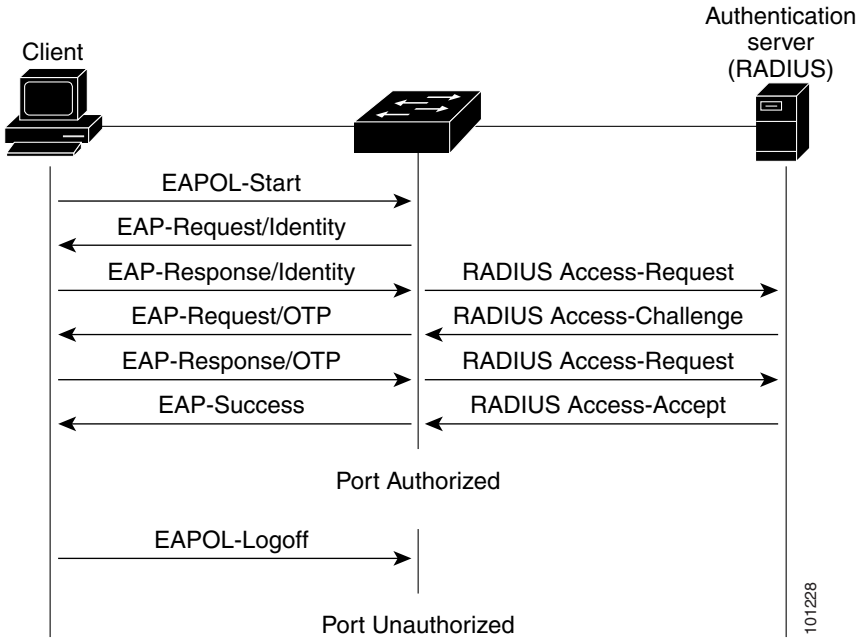

If 802.1x authentication times out while waiting for an EAPOL message exchange and MAC authentication bypass is enabled, the switch can authorize the client when the switch detects an Ethernet packet from the client. The switch uses the MAC address of the client as its identity and includes this information in the RADIUS-access/request frame that is sent to the RADIUS server. After the server sends the switch the RADIUS-access/accept frame (authorization is successful), the port becomes authorized. If authorization fails and a guest VLAN is specified, the switch assigns the port to the guest VLAN. If the switch detects an EAPOL packet while waiting for an Ethernet packet, the switch stops the MAC authentication bypass process and stops 802.1x authentication.

Figure 13-4 Message Exchange During MAC Authentication Bypass

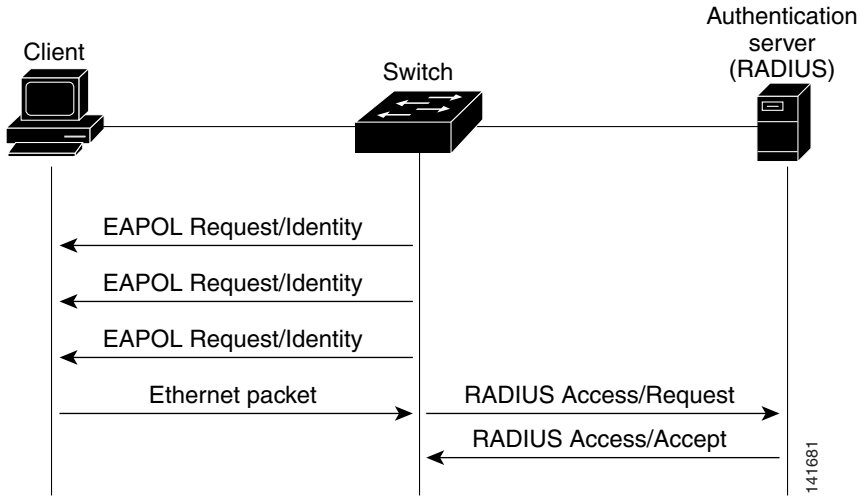

Authentication Manager

Port-Based Authentication Methods

Table 13-1 lists the authentication methods supported in these host modes:

- Single host—Only one data or voice host (client) can be authenticated on a port.
- Multiple host—Multiple data hosts can be authenticated on the same port. (If a port becomes unauthorized in multiple-host mode, the switch denies network access to all of the attached clients.)
- Multidomain authentication (MDA)—Both a data device and voice device can be authenticated on the same switch port. The port is divided into a data domain and a voice domain.
- Multiple authentication—Multiple hosts can authenticate on the data VLAN. This mode also allows one client on the VLAN if a voice VLAN is configured.

Table 13-1 802.1x Features

Authentication Method	Mode			
	Single Host	Multiple Host	MDA[1]	Multiple Authentication[2]
802.1x	VLAN assignment Per-user ACL Filter-ID attribute Downloadable ACL[3] Redirect URL [3]	VLAN assignment Per-user ACL Filter-ID attribute Downloadable ACL[4] Redirect URL [3]	VLAN assignment Per-user ACL[3] Filter-Id attribute[3] Downloadable ACL[3] Redirect URL[3]	Per-user ACL[3] Filter-Id attribute[3] Downloadable ACL[3] Redirect URL[3]
MAC authentication bypass	VLAN assignment Per-user ACL Filter-ID attribute Downloadable ACL[3] Redirect URL[3]	VLAN assignment Per-user ACL Filter-ID attribute Downloadable ACL[3] Redirect URL[3]	VLAN assignment Per-user ACL[3] Filter-Id attribute[3] Downloadable ACL[3] Redirect URL[3]	Per-user ACL[3] Filter-Id attribute[3] Downloadable ACL[3] Redirect URL[3]
Standalone web authentication[4]	Proxy ACL, Filter-Id attribute, downloadable ACL[2]			
NAC Layer 2 IP validation	Filter-Id attribute[3] Downloadable ACL Redirect URL	Filter-Id attribute[3] Downloadable ACL Redirect URL	Filter-Id attribute[3] Downloadable ACL Redirect URL	Filter-Id attribute[3] Downloadable ACL[3] Redirect URL[3]
Web authentication as fallback method[5]	Proxy ACL Filter-Id attribute[3] Downloadable ACL[3]	Proxy ACL Filter-Id attribute[3] Downloadable ACL[3]	Proxy ACL Filter-Id attribute[3] Downloadable ACL[3]	Proxy ACL[3] Filter-Id attribute[3] Downloadable ACL[3]

1. MDA = Multidomain authentication.
2. Also referred to as *multiauth*.
3. Supported in Cisco IOS Release 12.2(50)SE and later.
4. Supported in Cisco IOS Release 12.2(50)SE and later.
5. For clients that do not support 802.1x authentication.

Per-User ACLs and Filter-Ids

In releases earlier than Cisco IOS Release 12.2(50)SE, per-user ACLs and filter IDs were only supported in single-host mode. In Cisco IOS Release 12.2(50), support was added for MDA- and multiauth-enabled ports. In 12.2(52)SE and later, support was added for ports in multihost mode.

In releases earlier than Cisco IOS Release 12.2(50)SE, an ACL configured on the switch is not compatible with an ACL configured on another device running Cisco IOS software, such as a Catalyst 6500 switch.

In Cisco IOS Release 12.2(50)SE or later, the ACLs configured on the switch are compatible with other devices running the Cisco IOS release.

Note You can only set **any** as the source in the ACL.

Note For any ACL configured for multiple-host mode, the source portion of statement must be *any.* (For example, **permit icmp** *any* **host 10.10.1.1**.)

You must specify *any* in the source ports of any defined ACL. Otherwise, the ACL cannot be applied and authorization fails. Single host is the only exception to support backward compatibility.

More than one host can be authenticated on MDA- enabled and multiauth ports. The ACL policy applied for one host does not effect the traffic of another host.

If only one host is authenticated on a multihost port, and the other hosts gain network access without authentication, the ACL policy for the first host can be applied to the other connected hosts by specifying *any* in the source address.

Authentication Manager CLI Commands

The authentication-manager interface-configuration commands control all the authentication methods, such as 802.1x, MAC authentication bypass, and web authentication. The authentication manager commands determine the priority and order of authentication methods applied to a connected host.

The authentication manager commands control generic authentication features, such as host-mode, violation mode, and the authentication timer. Generic authentication commands include the **authentication host-mode**, **authentication violation**, and **authentication timer** interface configuration commands.

802.1x-specific commands begin with the **dot1x** or **authentication** keyword. For example, the **authentication port-control auto** interface configuration command enables authentication on an interface. However, the **dot1x system-authentication control g**lobal configuration command only globally enables or disables 802.1x authentication.

Note If 802.1x authentication is globally disabled, other authentication methods are still enabled on that port, such as web authentication.

You can filter out verbose system messages generated by the authentication manager. The filtered content typically relates to authentication success. You can also filter verbose messages for 802.1x authentication and MAB authentication. There is a separate command for each authentication method:

- The **no authentication logging verbose** global configuration command filters verbose messages from the authentication manager.
- The **no dot1x logging verbose** global configuration command filters 802.1x authentication verbose messages.
- The **no mab logging verbose** global configuration command filters MAC authentication bypass (MAB) verbose messages

For more information, see the command reference for this release.

Ports in Authorized and Unauthorized States

During 802.1x authentication, depending on the switch port state, the switch can grant a client access to the network. The port starts in the *unauthorized* state. While in this state, the port that is not configured as a voice VLAN port disallows all ingress and egress traffic except for 802.1x authentication, CDP, and STP packets. When a client is successfully authenticated, the port changes to the *authorized* state, allowing all traffic for the client to flow normally. If the port is configured as a voice VLAN port, the port allows VoIP traffic and 802.1x protocol packets before the client is successfully authenticated.

If a client that does not support 802.1x authentication connects to an unauthorized 802.1x port, the switch requests the client's identity. In this situation, the client does not respond to the request, the port remains in the unauthorized state, and the client is not granted access to the network.

In contrast, when an 802.1x-enabled client connects to a port that is not running the 802.1x standard, the client initiates the authentication process by sending the EAPOL-start frame. When no response is received, the client sends the request for a fixed number of times. Because no response is received, the client begins sending frames as if the port is in the authorized state.

You control the port authorization state by using the **authentication port-control** interface configuration command and these keywords:

- **force-authorized**—Disables 802.1x authentication and causes the port to change to the authorized state without any authentication exchange required. The port sends and receives normal traffic without 802.1x-based authentication of the client. This is the default setting.
- **force-unauthorized**—Causes the port to remain in the unauthorized state, ignoring all attempts by the client to authenticate. The switch cannot provide authentication services to the client through the port.
- **auto**—Enables 802.1x authentication and causes the port to begin in the unauthorized state, allowing only EAPOL frames to be sent and received through the port. The authentication process begins when the link state of the port changes from down to up or when an EAPOL-start frame is received. The switch requests the identity of the client and begins relaying authentication messages between the client and the authentication server. Each client attempting to access the network is uniquely identified by the switch by using the client MAC address.

If the client is successfully authenticated (receives an Accept frame from the authentication server), the port state changes to authorized, and all frames from the authenticated client are allowed through the port. If the authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the switch can resend the request. If no response is received from the server after the specified number of attempts, authentication fails, and network access is not granted.

When a client logs off, it sends an EAPOL-logoff message, causing the switch port to change to the unauthorized state.

If the link state of a port changes from up to down, or if an EAPOL-logoff frame is received, the port returns to the unauthorized state.

802.1x Host Mode

You can configure an 802.1x port for single-host or for multiple-hosts mode. In single-host mode (see Figure 13-1 on page 13-2), only one client can be connected to the 802.1x-enabled switch port. The switch detects the client by sending an EAPOL frame when the port link state changes to the up state. If a client leaves or is replaced with another client, the switch changes the port link state to down, and the port returns to the unauthorized state.

In multiple-hosts mode, you can attach multiple hosts to a single 802.1x-enabled port. Figure 13-5 on page 13-10 shows 802.1x port-based authentication in a wireless LAN. In this mode, only one of the attached clients must be authorized for all clients to be granted network access. If the port becomes unauthorized (reauthentication fails or an EAPOL-logoff message is received), the switch denies network access to all of the attached clients. In this topology, the wireless access point is responsible for authenticating the clients attached to it, and it also acts as a client to the switch.

Figure 13-5 Multiple Host Mode Example

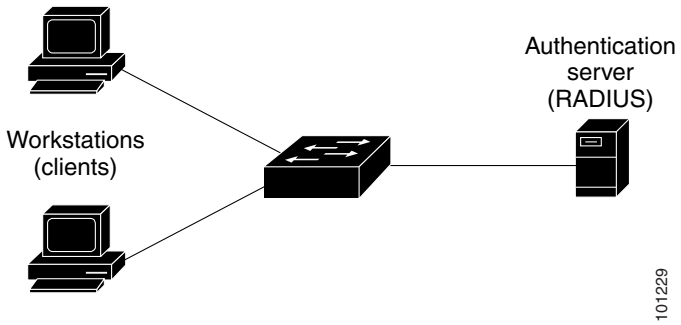

The switch supports multidomain authentication (MDA), which allows both a data device and a voice device, such as an IP Phone (Cisco or non-Cisco), to connect to the same switch port. For more information, see the "Multidomain Authentication" section on page 13-10.

Multidomain Authentication

The switch supports multidomain authentication (MDA), which allows both a data device and voice device, such as an IP phone (Cisco or non-Cisco), to authenticate on the same switch port. The port is divided into a data domain and a voice domain.

MDA does not enforce the order of device authentication. However, for best results, we recommend that a voice device is authenticated before a data device on an MDA-enabled port.

Follow these guidelines for configuring MDA:

- To configure a switch port for MDA, see the "Configuring the Host Mode" section on page 13-38.
- You must configure the voice VLAN for the IP phone when the host mode is set to multidomain. For more information, see Chapter 17, "Configuring VLANs."
- To authorize a voice device, the AAA server must be configured to send a Cisco Attribute-Value (AV) pair attribute with a value of `device-traffic-class=voice`. Without this value, the switch treats the voice device as a data device.
- The guest VLAN and restricted VLAN features only apply to the data devices on an MDA-enabled port. The switch treats a voice device that fails authorization as a data device.
- If more than one device attempts authorization on either the voice or the data domain of a port, it is error disabled.
- Until a device is authorized, the port drops its traffic. Non-Cisco IP phones or voice devices are allowed into both the data and voice VLANs. The data VLAN allows the voice device to contact a DHCP server to obtain an IP address and acquire the voice VLAN information. After the voice device starts sending on the voice VLAN, its access to the data VLAN is blocked.

- A voice device MAC address that is binding on the data VLAN is not counted towards the port security MAC address limit.
- MDA can use MAC authentication bypass as a fallback mechanism to allow the switch port to connect to devices that do not support 802.1x authentication. For more information, see the "MAC Authentication Bypass Guidelines" section on page 13-33.
- When a *data* or a *voice* device is detected on a port, its MAC address is blocked until authorization succeeds. If the authorization fails, the MAC address remains blocked for 5 minutes.
- If more than five devices are detected on the *data* VLAN or more than one voice device is detected on the *voice* VLAN while a port is unauthorized, the port is error disabled.
- When a port host mode changes from single- or multihost to multidomain mode, an authorized data device remains authorized on the port. However, a Cisco IP phone on the port voice VLAN is automatically removed and must be reauthenticated on that port.
- Active fallback mechanisms such as guest VLAN and restricted VLAN remain configured after a port changes from single-host or multiple-host mode to multidomain mode.
- Switching a port host mode from multidomain to single-host or multiple-hosts mode removes all authorized devices from the port.
- If a data domain is authorized first and placed in the guest VLAN, non-802.1x-capable voice devices need their packets tagged on the voice VLAN to trigger authentication. The phone need not need to send tagged traffic. (The same is true for an 802.1x-capable phone.)
- We do not recommend per-user ACLs with an MDA-enabled port. An authorized device with a per-user ACL policy might impact traffic on both the port voice and data VLANs. You can use only one device on the port to enforce per-user ACLs.

For more information, see the "Configuring the Host Mode" section on page 13-38.

802.1x Multiple Authentication Mode

Multiple-authentication (multiauth) mode allows multiple authenticated clients on the data VLAN. Each host is individually authenticated. If a voice VLAN is configured, this mode also allows one client on the VLAN. (If the port detects any additional voice clients, they are discarded from the port, but no violation errors occur.)

If a hub or access point is connected to an 802.1x-enabled port, each connected client must be authenticated.

For non-802.1x devices, you can use MAC authentication bypass or web authentication as the per-host authentication fallback method to authenticate different hosts with different methods on a single port.

There is no limit to the number of data hosts can authenticate on a multiauthport. However, only one voice device is allowed if the voice VLAN is configured. Since there is no host limit defined violation will not be trigger, if a second voice is seen we silently discard it but do not trigger violation.

For MDA functionality on the voice VLAN, multiple-authentication mode assigns authenticated devices to either a data or a voice VLAN, depending on the VSAs received from the authentication server.

Note When a port is in multiple-authentication mode, the guest VLAN and the authentication-failed VLAN features do not activate.

For more information about critical authentication mode and the critical VLAN, see the "802.1x Authentication with Inaccessible Authentication Bypass" section on page 13-22.

For more information about configuring multiauth mode on a port, see the "Configuring the Host Mode" section on page 13-38.

MAC Move

When a MAC address is authenticated on one switch port, that address is not allowed on another authentication manager-enabled port of the switch. If the switch detects that same MAC address on another authentication manager-enabled port, the address is not allowed.

There are situations where a MAC address might need to move from one port to another on the same switch. For example, when there is another device (for example a hub or an IP phone) between an authenticated host and a switch port, you might want to disconnect the host from the device and connect it directly to another port on the same switch.

You can globally enable MAC move so the device is reauthenticated on the new port. When a host moves to a second port, the session on the first port is deleted, and the host is reauthenticated on the new port.

MAC move is supported on all host modes. (The authenticated host can move to any port on the switch, no matter which host mode is enabled on the that port.)

When a MAC address moves from one port to another, the switch terminates the authenticated session on the original port and initiates a new authentication sequence on the new port.

The MAC move feature applies to both voice and data hosts.

Note In open authentication mode, a MAC address is immediately moved from the original port to the new port, with no requirement for authorization on the new port.

For more information see the "Configuring Optional 802.1x Authentication Features" section on page 13-40.

MAC Replace

The MAC replace feature can be configured to address the violation that occurs when a host attempts to connect to a port where another host was previously authenticated.

Note This feature does not apply to ports in multiauth mode, because violations are not triggered in that mode. It does not apply to ports in multiple host mode, because in that mode, only the first host requires authentication.

If you configure the **authentication violation** interface configuration command with the **replace** keyword, the authentication process on a port in multidomain mode is:

- A new MAC address is received on a port with an existing authenticated MAC address.
- The authentication manager replaces the MAC address of the current data host on the port with the new MAC address.
- The authentication manager initiates the authentication process for the new MAC address.
- If the authentication manager determines that the new host is a voice host, the original voice host is removed.

If a port is in open authentication mode, any new MAC address is immediately added to the MAC address table.

For more information see the "Configuring Optional 802.1x Authentication Features" section on page 13-40.

802.1x Accounting

The 802.1x standard defines how users are authorized and authenticated for network access but does not keep track of network usage. 802.1x accounting is disabled by default. You can enable 802.1x accounting to monitor this activity on 802.1x-enabled ports:

- User successfully authenticates.
- User logs off.
- Link-down occurs.
- Reauthentication successfully occurs.
- Reauthentication fails.

The switch does not log 802.1x accounting information. Instead, it sends this information to the RADIUS server, which must be configured to log accounting messages.

802.1x Accounting Attribute-Value Pairs

The information sent to the RADIUS server is represented in the form of Attribute-Value (AV) pairs. These AV pairs provide data for different applications. (For example, a billing application might require information that is in the Acct-Input-Octets or the Acct-Output-Octets attributes of a RADIUS packet.)

AV pairs are automatically sent by a switch that is configured for 802.1x accounting. Three types of RADIUS accounting packets are sent by a switch:

- START—Sent when a new user session starts
- INTERIM—Sent during an existing session for updates
- STOP—Sent when a session terminates

Table 13-2 Accounting AV Pairs

Attribute Number	AV Pair Name	START	INTERIM	STOP
Attribute[1]	User-Name	Always	Always	Always
Attribute[4]	NAS-IP-Address	Always	Always	Always
Attribute[5]	NAS-Port	Always	Always	Always
Attribute[8]	Framed-IP-Address	Never	Sometimes[1]	Sometimes[1]
Attribute[25]	Class	Always	Always	Always
Attribute[30]	Called-Station-ID	Always	Always	Always
Attribute[31]	Calling-Station-ID	Always	Always	Always
Attribute[40]	Acct-Status-Type	Always	Always	Always
Attribute[41]	Acct-Delay-Time	Always	Always	Always
Attribute[42]	Acct-Input-Octets	Never	Always	Always

Table 13-2 Accounting AV Pairs (continued)

Attribute Number	AV Pair Name	START	INTERIM	STOP
Attribute[43]	Acct-Output-Octets	Never	Always	Always
Attribute[44]	Acct-Session-ID	Always	Always	Always
Attribute[45]	Acct-Authentic	Always	Always	Always
Attribute[46]	Acct-Session-Time	Never	Always	Always
Attribute[49]	Acct-Terminate-Cause	Never	Never	Always
Attribute[61]	NAS-Port-Type	Always	Always	Always

1. The Framed-IP-Address AV pair is sent only if a valid Dynamic Host Control Protocol (DHCP) binding exists for the host in the DHCP snooping bindings table.

You can view the AV pairs that are being sent by the switch by entering the **debug radius accounting** privileged EXEC command. For more information about this command, see the *Cisco IOS Debug Command Reference, Release 12.2*.

For more information about AV pairs, see RFC 3580, "802.1x Remote Authentication Dial In User Service (RADIUS) Usage Guidelines."

802.1x Readiness Check

The 802.1x readiness check monitors 802.1x activity on all the switch ports and displays information about the devices connected to the ports that support 802.1x. You can use this feature to determine if the devices connected to the switch ports are 802.1x-capable. You use an alternate authentication such as MAC authentication bypass or web authentication for the devices that do not support 802.1x functionality.

This feature only works if the supplicant on the client supports a query with the NOTIFY EAP notification packet. The client must respond within the 802.1x timeout value.

Follow these guidelines to enable the readiness check on the switch:

- The readiness check is typically used before 802.1x is enabled on the switch.
- The 802.1x readiness check is allowed on all ports that can be configured for 802.1x. The readiness check is not available on a port that is configured as **dot1x force-unauthorized**.
- If you use the **dot1x test eapol-capable** privileged EXEC command without specifying an interface, all the ports on the switch stack are tested.
- When you configure the **dot1x test eapol-capable** command on an 802.1x-enabled port, and the link comes up, the port queries the connected client about its 802.1x capability. When the client responds with a notification packet, it is 802.1x-capable. A syslog message is generated if the client responds within the timeout period. If the client does not respond to the query, the client is not 802.1x-capable. No syslog message is generated.
- The readiness check can be sent on a port that handles multiple hosts (for example, a PC that is connected to an IP phone). A syslog message is generated for each of the clients that respond to the readiness check within the timer period.

For information on configuring the switch for the 802.1x readiness check, see the "Configuring 802.1x Readiness Check" section on page 13-36.

802.1x Authentication with VLAN Assignment

The RADIUS server sends the VLAN assignment to configure the switch port. The RADIUS server database maintains the username-to-VLAN mappings, assigning the VLAN based on the username of the client connected to the switch port. You can use this feature to limit network access for certain users.

When a voice device is authorized and the RADIUS server returns an authorized VLAN, the voice VLAN on the port is configured to send and receive packets on the assigned voice VLAN. Voice VLAN assignment behaves the same as data VLAN assignment on multidomain authentication (MDA)-enabled ports. For more information, see the "Multidomain Authentication" section on page 13-10.

When configured on the switch and the RADIUS server, 802.1x authentication with VLAN assignment has these characteristics:

- If no VLAN is supplied by the RADIUS server or if 802.1x authentication is disabled, the port is configured in its access VLAN after successful authentication. Recall that an access VLAN is a VLAN assigned to an access port. All packets sent from or received on this port belong to this VLAN.
- If 802.1x authentication is enabled but the VLAN information from the RADIUS server is not valid, authorization fails and configured VLAN remains in use. This prevents ports from appearing unexpectedly in an inappropriate VLAN because of a configuration error.

  Configuration errors could include specifying a VLAN for a routed port, a malformed VLAN ID, a nonexistent or internal (routed port) VLAN ID, an RSPAN VLAN, a shut down or suspended VLAN. In the case of a mutlidomain host port, configuration errors can also be due to an attempted assignment of a data VLAN that matches the configured or assigned voice VLAN ID (or the reverse).
- If 802.1x authentication is enabled and all information from the RADIUS server is valid, the authorized device is placed in the specified VLAN after authentication.
- If the multiple-hosts mode is enabled on an 802.1x port, all hosts are placed in the same VLAN (specified by the RADIUS server) as the first authenticated host.
- Enabling port security does not impact the RADIUS server-assigned VLAN behavior.
- If 802.1x authentication is disabled on the port, it is returned to the configured access VLAN and configured voice VLAN.
- If an 802.1x port is authenticated and put in the RADIUS server-assigned VLAN, any change to the port access VLAN configuration does not take effect. In the case of a multidomain host, the same applies to voice devices when the port is fully authorized with these exceptions:
  - If the VLAN configuration change of one device results in matching the other device configured or assigned VLAN, then authorization of all devices on the port is terminated and multidomain host mode is disabled until a valid configuration is restored where data and voice device configured VLANs no longer match.
  - If a voice device is authorized and is using a downloaded voice VLAN, the removal of the voice VLAN configuration, or modifying the configuration value to *dot1p* or *untagged* results in voice device un-authorization and the disablement of multi-domain host mode.

When the port is in the force authorized, force unauthorized, unauthorized, or shutdown state, it is put into the configured access VLAN.

The 802.1x authentication with VLAN assignment feature is not supported on trunk ports, dynamic ports, or with dynamic-access port assignment through a VLAN Membership Policy Server (VMPS).

To configure VLAN assignment you need to perform these tasks:

- Enable AAA authorization by using the **network** keyword to allow interface configuration from the RADIUS server.
- Enable 802.1x authentication. (The VLAN assignment feature is automatically enabled when you configure 802.1x authentication on an access port.)
- Assign vendor-specific tunnel attributes in the RADIUS server. The RADIUS server must return these attributes to the switch:
  - [64] Tunnel-Type = VLAN
  - [65] Tunnel-Medium-Type = 802
  - [81] Tunnel-Private-Group-ID = VLAN name, VLAN ID, or VLAN-Group
  - [83] Tunnel-Preference

  Attribute [64] must contain the value *VLAN* (type 13). Attribute [65] must contain the value *802* (type 6). Attribute [81] specifies the *VLAN name* or *VLAN ID* assigned to the 802.1x-authenticated user.

For examples of tunnel attributes, see the "Configuring Vendor-Specific RADIUS Attributes: Examples" section on page 12-46.

Voice Aware 802.1x Security

You use the voice aware 802.1x security feature on the switch to disable only the VLAN on which a security violation occurs, whether it is a data or voice VLAN. You can use this feature in IP phone deployments where a PC is connected to the IP phone. A security violation found on the data VLAN results in the shutdown of only the data VLAN. The traffic on the voice VLAN flows through the switch without interruption.

Follow these guidelines to configure voice aware 802.1x voice security on the switch:

- You enable voice aware 802.1x security by entering the **errdisable detect cause security-violation shutdown vlan** global configuration command. You disable voice aware 802.1x security by entering the **no** version of this command. This command applies to all 802.1x-configured ports in the switch.

Note If you do not include the **shutdown vlan** keywords, the entire port is shut down when it enters the error-disabled state.

- If you use the **errdisable recovery cause security-violation** global configuration command to configure error-disabled recovery, the port is automatically reenabled. If error-disabled recovery is not configured for the port, you reenable it by using the **shutdown** and **no-shutdown** interface configuration commands.
- You can reenable individual VLANs by using the **clear errdisable interface** *interface-id* **vlan** [*vlan-list*] privileged EXEC command. If you do not specify a range, all VLANs on the port are enabled.

802.1x Authentication with Per-User ACLs

You can enable per-user access control lists (ACLs) to provide different levels of network access and service to an 802.1x-authenticated user. When the RADIUS server authenticates a user connected to an 802.1x port, it retrieves the ACL attributes based on the user identity and sends them to the switch. The switch applies the attributes to the 802.1x port for the duration of the user session. The switch removes the per-user ACL configuration when the session is over, if authentication fails, or if a link-down condition occurs. The switch does not save RADIUS-specified ACLs in the running configuration. When the port is unauthorized, the switch removes the ACL from the port.

You can configure router ACLs and input port ACLs on the same switch. However, a port ACL takes precedence over a router ACL. If you apply input port ACL to an interface that belongs to a VLAN, the port ACL takes precedence over an input router ACL applied to the VLAN interface. Incoming packets received on the port to which a port ACL is applied are filtered by the port ACL. Incoming routed packets received on other ports are filtered by the router ACL. Outgoing routed packets are filtered by the router ACL. To avoid configuration conflicts, you should carefully plan the user profiles stored on the RADIUS server.

RADIUS supports per-user attributes, including vendor-specific attributes. These vendor-specific attributes (VSAs) are in octet-string format and are passed to the switch during the authentication process. The VSAs used for per-user ACLs are `inacl#<n>` for the ingress direction and `outacl#<n>` for the egress direction. MAC ACLs are supported only in the ingress direction. The switch supports VSAs only in the ingress direction. It does not support port ACLs in the egress direction on Layer 2 ports. For more information, see Chapter 37, “Configuring Network Security with ACLs.”

Use only the extended ACL syntax style to define the per-user configuration stored on the RADIUS server. When the definitions are passed from the RADIUS server, they are created by using the extended naming convention. However, if you use the Filter-Id attribute, it can point to a standard ACL.

You can use the Filter-Id attribute to specify an inbound or outbound ACL that is already configured on the switch. The attribute contains the ACL number followed by *.in* for ingress filtering or *.out* for egress filtering. If the RADIUS server does not allow the *.in* or *.out* syntax, the access list is applied to the outbound ACL by default. Because of limited support of Cisco IOS access lists on the switch, the Filter-Id attribute is supported only for IP ACLs numbered 1 to 199 and 1300 to 2699 (IP standard and IP extended ACLs).

The maximum size of the per-user ACL is 4000 ASCII characters but is limited by the maximum size of RADIUS-server per-user ACLs.

For examples of vendor-specific attributes, see the “Configuring Vendor-Specific RADIUS Attributes: Examples” section on page 12-46. For more information about configuring ACLs, see Chapter 37, “Configuring Network Security with ACLs.”

Note Per-user ACLs are supported only in single-host mode.

To configure per-user ACLs, you need to perform these tasks:

- Enable AAA authentication.
- Enable AAA authorization by using the **network** keyword to allow interface configuration from the RADIUS server.
- Enable 802.1x authentication.
- Configure the user profile and VSAs on the RADIUS server.
- Configure the 802.1x port for single-host mode.

For more configuration information, see the “Authentication Manager” section on page 13-6.

802.1x Authentication with Downloadable ACLs and Redirect URLs

You can download ACLs and redirect URLs from a RADIUS server to the switch during 802.1x authentication or MAC authentication bypass of the host. You can also download ACLs during web authentication.

Note A downloadable ACL is also referred to as a *dACL*.

If more than one host is authenticated and the host is in single-host, MDA, or multiple-authentication mode, the switch changes the source address of the ACL to the host IP address.

You can apply the ACLs and redirect URLs to all the devices connected to the 802.1x-enabled port.

If no ACLs are downloaded during 802.1x authentication, the switch applies the static default ACL on the port to the host. On a voice VLAN port configured in multi-auth or MDA mode, the switch applies the ACL only to the phone as part of the authorization policies.

Note The auth-default ACL does not appear in the running configuration.

The auth-default ACL is created when at least one host with an authorization policy is detected on the port. The auth-default ACL is removed from the port when the last authenticated session ends. You can configure the auth-default ACL by using the **ip access-list extended auth-default-acl** global configuration command.

Note The auth-default ACL does not support Cisco Discovery Protocol (CDP) bypass in the single host mode. You must configure a static ACL on the interface to support CDP bypass.

The 802.1x and MAB authentication methods support two authentication modes, *open* and *closed*. If there is no static ACL on a port in *closed* authentication mode:

- An auth-default-ACL is created.
- The auth-default-ACL allows only DHCP traffic until policies are enforced.
- When the first host authenticates, the authorization policy is applied without IP address insertion.
- When a second host is detected, the policies for the first host are refreshed, and policies for the first and subsequent sessions are enforced with IP address insertion.

If there is no static ACL on a port in *open* authentication mode:

- An auth-default-ACL-OPEN is created and allows all traffic.
- Policies are enforced with IP address insertion to prevent security breaches.
- Web authentication is subject to the auth-default-ACL-OPEN.

To control access for hosts with no authorization policy, you can configure a directive. The supported values for the directive are *open* and *default*. When you configure the *open* directive, all traffic is allowed. The *default* directive subjects traffic to the access provided by the port. You can configure the directive either in the user profile on the AAA server or on the switch. To configure the directive on the AAA server, use the **authz-directive =** *open*/*default* global command. To configure the directive on the switch, use the **epm access-control open** global configuration command.

Note The default value of the directive is *default*.

If a host falls back to web authentication on a port without a configured ACL:

- If the port is in open authentication mode, the auth-default-ACL-OPEN is created.
- If the port is in closed authentication mode, the auth-default-ACL is created.

The access control entries (ACEs) in the fallback ACL are converted to per-user entries. If the configured fallback profile does not include a fallback ACL, the host is subject to the auth-default-ACL associated with the port.

Note If you use a custom logo with web authentication and it is stored on an external server, the port ACL must allow access to the external server before authentication. You must either configure a static port ACL or change the auth-default-ACL to provide appropriate access to the external server.

Cisco Secure ACS and Attribute-Value Pairs for the Redirect URL

The switch uses these *cisco-av-pair* VSAs:

- url-redirect is the HTTP to HTTPS URL.
- url-redirect-acl is the switch ACL name or number.

The switch uses the CiscoSecure-Defined-ACL attribute value pair to intercept an HTTP or HTTPS request from the end point device. The switch then forwards the client web browser to the specified redirect address. The url-redirect attribute value pair on the Cisco Secure ACS contains the URL to which the web browser is redirected. The url-redirect-acl attribute value pair contains the name or number of an ACL that specifies the HTTP or HTTPS traffic to redirect. Traffic that matches a permit ACE in the ACL is redirected.

Note Define the URL redirect ACL and the default port ACL on the switch.

If a redirect URL is configured for a client on the authentication server, a default port ACL on the connected client switch port must also be configured.

Cisco Secure ACS and Attribute-Value Pairs for Downloadable ACLs

You can set the CiscoSecure-Defined-ACL Attribute-Value pair on the Cisco Secure ACS with the RADIUS cisco-av-pair vendor-specific attributes (VSAs). This pair specifies the names of the downloadable ACLs on the Cisco Secure ACS with the #ACL#-IP-*name-number* attribute.

- The *name* is the ACL name.
- The *number* is the version number (for example, 3f783768).

If a downloadable ACL is configured for a client on the authentication server, a default port ACL on the connected client switch port must also be configured.

If the default ACL is configured on the switch and the Cisco Secure ACS sends a host-access-policy to the switch, it applies the policy to traffic from the host connected to a switch port. If the policy does not apply, the switch applies the default ACL. If the Cisco Secure ACS sends the switch a downloadable ACL, this ACL takes precedence over the default ACL that is configured on the switch port. However, if the switch receives an host access policy from the Cisco Secure ACS but the default ACL is not configured, the authorization failure is declared.

For configuration details, see the "Authentication Manager" section on page 13-6 and the "Configuring 802.1x Authentication with Downloadable ACLs and Redirect URLs" section on page 13-48.

VLAN ID-Based MAC Authentication

You can use VLAN ID-based MAC authentication if you want to authenticate hosts based on a static VLAN ID instead of a downloadable VLAN. When you have a static VLAN policy configured on your switch, VLAN information is sent to an IAS (Microsoft) RADIUS server along with the MAC address of each host for authentication. The VLAN ID configured on the connected port is used for MAC authentication. By using VLAN ID-based MAC authentication with an IAS server, you can have a fixed number of VLANs in the network.

The feature also limits the number of VLANs monitored and handled by STP. The network can be managed as a fixed VLAN.

Note This feature is not supported on Cisco ACS Server. (The ACS server ignores the sent VLAN-IDs for new hosts and only authenticates based on the MAC address.)

For configuration information, see the "Configuring Optional 802.1x Authentication Features" section on page 13-40. Additional configuration is similar MAC authentication bypass, as described in the "Configuring 802.1x User Distribution" section on page 13-46.

802.1x Authentication with Guest VLAN

You can configure a guest VLAN for each 802.1x port on the switch to provide limited services to clients, such as downloading the 802.1x client. These clients might be upgrading their system for 802.1x authentication, and some hosts, such as Windows 98 systems, might not be 802.1x-capable.

When you enable a guest VLAN on an 802.1x port, the switch assigns clients to a guest VLAN when the switch does not receive a response to its EAP request/identity frame or when EAPOL packets are not sent by the client. The port is automatically set to multi-host mode.

The switch maintains the EAPOL packet history. If an EAPOL packet is detected on the interface during the lifetime of the link, the switch determines that the device connected to that interface is an 802.1x-capable supplicant, and the interface does not change to the guest VLAN state. EAPOL history is cleared if the interface link status goes down. If no EAPOL packet is detected on the interface, the interface changes to the guest VLAN state.

If devices send EAPOL packets to the switch during the lifetime of the link, the switch no longer allows clients that fail authentication access to the guest VLAN.

If the switch is trying to authorize an 802.1x-capable voice device and the AAA server is unavailable, the authorization attempt fails, but the detection of the EAPOL packet is saved in the EAPOL history. When the AAA server becomes available, the switch authorizes the voice device. However, the switch no longer allows other devices access to the guest VLAN. To prevent this situation, use one of these command sequences:

- Enter the **authentication event no-response action authorize vlan** *vlan-id* interface configuration command to allow access to the guest VLAN.
- Enter the **shutdown** interface configuration command followed by the **no shutdown** interface configuration command to restart the port.

Note If an EAPOL packet is detected after the interface has changed to the guest VLAN, the interface reverts to an unauthorized state, and 802.1x authentication restarts.

Any number of 802.1x-incapable clients are allowed access when the switch port is moved to the guest VLAN. If an 802.1x-capable client joins the same port on which the guest VLAN is configured, the port is put into the unauthorized state in the user-configured access VLAN, and authentication is restarted.

Guest VLANs are supported on 802.1x ports in single host, multiple host, or multi-domain modes.

You can configure any active VLAN except an RSPAN VLAN, a private VLAN, or a voice VLAN as an 802.1x guest VLAN. The guest VLAN feature is not supported on internal VLANs (routed ports) or trunk ports; it is supported only on access ports.

The switch supports *MAC authentication bypass*. When MAC authentication bypass is enabled on an 802.1x port, the switch can authorize clients based on the client MAC address when 802.1x authentication times out while waiting for an EAPOL message exchange. After detecting a client on an 802.1x port, the switch waits for an Ethernet packet from the client. The switch sends the authentication server a RADIUS-access/request frame with a username and password based on the MAC address. If authorization succeeds, the switch grants the client access to the network. If authorization fails, the switch assigns the port to the guest VLAN if one is specified. For more information, see the "802.1x Authentication with MAC Authentication Bypass" section on page 13-25.

For more information, see the "Configuring a Guest VLAN" section on page 13-42.

802.1x Authentication with Restricted VLAN

You can configure a restricted VLAN (also referred to as an *authentication failed VLAN*) for each 802.1x port on a switch to provide limited services to clients that cannot access the guest VLAN. These clients are 802.1x-compliant and cannot access another VLAN because they fail the authentication process. A restricted VLAN allows users without valid credentials in an authentication server (typically, visitors to an enterprise) to access a limited set of services. The administrator can control the services available to the restricted VLAN.

Note You can configure a VLAN to be both the guest VLAN and the restricted VLAN if you want to provide the same services to both types of users.

Without this feature, the client attempts and fails authentication indefinitely, and the switch port remains in the spanning-tree blocking state. With this feature, you can configure the switch port to be in the restricted VLAN after a specified number of authentication attempts (the default value is 3 attempts).

The authenticator counts the failed authentication attempts for the client. When this count exceeds the configured maximum number of authentication attempts, the port moves to the restricted VLAN. The failed attempt count increments when the RADIUS server replies with either an *EAP failure* or an empty response without an EAP packet. When the port moves into the restricted VLAN, the failed attempt counter resets.

Users who fail authentication remain in the restricted VLAN until the next reauthentication attempt. A port in the restricted VLAN tries to reauthenticate at configured intervals (the default is 60 seconds). If reauthentication fails, the port remains in the restricted VLAN. If reauthentication is successful, the port moves either to the configured VLAN or to a VLAN sent by the RADIUS server. You can disable reauthentication. If you do this, the only way to restart the authentication process is for the port to receive a *link dow*n or *EAP logoff* event. We recommend that you keep reauthentication enabled if a client might connect through a hub. When a client disconnects from the hub, the port might not receive the *link down* or *EAP logoff* event.

After a port moves to the restricted VLAN, a simulated EAP success message is sent to the client. This prevents clients from indefinitely attempting authentication. Some clients (for example, devices running Windows XP) cannot implement DHCP without EAP success.

Restricted VLANs are supported only on 802.1x ports in single-host mode and on Layer 2 ports.

You can configure any active VLAN except an RSPAN VLAN, a primary private VLAN, or a voice VLAN as an 802.1x restricted VLAN. The restricted VLAN feature is not supported on internal VLANs (routed ports) or trunk ports; it is supported only on access ports.

Other security features such as dynamic ARP inspection, DHCP snooping, and IP source guard can be configured independently on a restricted VLAN.

For more information, see the "Configuring a Restricted VLAN" section on page 13-43.

802.1x Authentication with Inaccessible Authentication Bypass

Use the inaccessible authentication bypass feature, also referred to as *critical authentication* or the *AAA fail policy,* when the switch cannot reach the configured RADIUS servers and new hosts cannot be authenticated. You can configure the switch to connect those hosts to *critical ports*.

When a new host tries to connect to the critical port, that host is moved to a user-specified access VLAN, the *critical VLAN*. The administrator gives limited authentication to the hosts.

When the switch tries to authenticate a host connected to a critical port, the switch checks the status of the configured RADIUS server. If a server is available, the switch can authenticate the host. However, if all the RADIUS servers are unavailable, the switch grants network access to the host and puts the port in the *critical-authentication* state, which is a special case of the authentication state.

Support on Multiple-Authentication Ports

When a port is configured on any host mode and the AAA server is unavailable, the port is then configured to multi-host mode and moved to the critical VLAN. To support this inaccessible bypass on multiple-authentication (multiauth) ports, use the **authentication event server dead action reinitialize vlan** *vlan-id* command. When a new host tries to connect to the critical port, that port is reinitialized and all the connected hosts are moved to the user-specified access VLAN.

This command is supported on all host modess.

Authentication Results

The behavior of the inaccessible authentication bypass feature depends on the authorization state of the port:

- If the port is unauthorized when a host connected to a critical port tries to authenticate and all servers are unavailable, the switch puts the port in the critical-authentication state in the RADIUS-configured or user-specified access VLAN.
- If the port is already authorized and reauthentication occurs, the switch puts the critical port in the critical-authentication state in the current VLAN, which might be the one previously assigned by the RADIUS server.
- If the RADIUS server becomes unavailable during an authentication exchange, the current exchange times out, and the switch puts the critical port in the critical-authentication state during the next authentication attempt.

You can configure the critical port to reinitialize hosts and move them out of the critical VLAN when the RADIUS server is again available. When this is configured, all critical ports in the critical-authentication state are automatically reauthenticated. For more information, see the command reference for this release and the "Configuring Inaccessible Authentication Bypass" section on page 13-44.

Feature Interactions

Inaccessible authentication bypass interacts with these features:

- Guest VLAN—Inaccessible authentication bypass is compatible with guest VLAN. When a guest VLAN is enabled on 8021.x port, the features interact as follows:
  - If at least one RADIUS server is available, the switch assigns a client to a guest VLAN when the switch does not receive a response to its EAP request/identity frame or when EAPOL packets are not sent by the client.
  - If all the RADIUS servers are not available and the client is connected to a critical port, the switch authenticates the client and puts the critical port in the critical-authentication state in the RADIUS-configured or user-specified access VLAN.
  - If all the RADIUS servers are not available and the client is not connected to a critical port, the switch might not assign clients to the guest VLAN if one is configured.
  - If all the RADIUS servers are not available and if a client is connected to a critical port and was previously assigned to a guest VLAN, the switch keeps the port in the guest VLAN.
- Restricted VLAN—If the port is already authorized in a restricted VLAN and the RADIUS servers are unavailable, the switch puts the critical port in the critical-authentication state in the restricted VLAN.
- 802.1x accounting—Accounting is not affected if the RADIUS servers are unavailable.
- Private VLAN—You can configure inaccessible authentication bypass on a private VLAN host port. The access VLAN must be a secondary private VLAN.
- Voice VLAN—Inaccessible authentication bypass is compatible with voice VLAN, but the RADIUS-configured or user-specified access VLAN and the voice VLAN must be different.
- Remote Switched Port Analyzer (RSPAN)—Do not configure an RSPAN VLAN as the RADIUS-configured or user-specified access VLAN for inaccessible authentication bypass.

802.1x Authentication with Voice VLAN Ports

A voice VLAN port is a special access port associated with two VLAN identifiers:

- VVID to carry voice traffic to and from the IP phone. The VVID is used to configure the IP phone connected to the port.
- PVID to carry the data traffic to and from the workstation connected to the switch through the IP phone. The PVID is the native VLAN of the port.

The IP phone uses the VVID for its voice traffic, regardless of the authorization state of the port. This allows the phone to work independently of 802.1x authentication.

In single-host mode, only the IP phone is allowed on the voice VLAN. In multiple-hosts mode, additional clients can send traffic on the voice VLAN after a supplicant is authenticated on the PVID. When multiple-hosts mode is enabled, the supplicant authentication affects both the PVID and the VVID.

A voice VLAN port becomes active when there is a link, and the device MAC address appears after the first CDP message from the IP phone. Cisco IP phones do not relay CDP messages from other devices. As a result, if several IP phones are connected in series, the switch recognizes only the one directly connected to it. When 802.1x authentication is enabled on a voice VLAN port, the switch drops packets from unrecognized IP phones more than one hop away.

When 802.1x authentication is enabled on a port, you cannot configure a port VLAN that is equal to a voice VLAN.

Note If you enable 802.1x authentication on an access port on which a voice VLAN is configured and to which a Cisco IP Phone is connected, the Cisco IP phone loses connectivity to the switch for up to 30 seconds.

For more information about voice VLANs, see Chapter 19, "Configuring Voice VLAN."

802.1x Authentication with Port Security

In general, Cisco does not recommend enabling port security when IEEE 802.1x is enabled. Since IEEE 802.1x enforces a single MAC address per port (or per VLAN when MDA is configured for IP telephony), port security is redundant and in some cases may interfere with expected IEEE 802.1x operations.

802.1x Authentication with Wake-on-LAN

The 802.1x authentication with the wake-on-LAN (WoL) feature allows dormant PCs to be powered when the switch receives a specific Ethernet frame, known as the *magic packet*. You can use this feature in environments where administrators need to connect to systems that have been powered down.

When a host that uses WoL is attached through an 802.1x port and the host powers off, the 802.1x port becomes unauthorized. The port can only receive and send EAPOL packets, and WoL magic packets cannot reach the host. When the PC is powered off, it is not authorized, and the switch port is not opened.

When the switch uses 802.1x authentication with WoL, the switch forwards traffic to unauthorized 802.1x ports, including magic packets. While the port is unauthorized, the switch continues to block ingress traffic other than EAPOL packets. The host can receive packets but cannot send packets to other devices in the network.

Note If PortFast is not enabled on the port, the port is forced to the bidirectional state.

When you configure a port as unidirectional by using the **authentication control-direction in** interface configuration command, the port changes to the spanning-tree forwarding state. The port can send packets to the host but cannot receive packets from the host.

When you configure a port as bidirectional by using the **authentication control-direction both** interface configuration command, the port is access-controlled in both directions. The port does not receive packets from or send packets to the host.

802.1x Authentication with MAC Authentication Bypass

You can configure the switch to authorize clients based on the client MAC address (see Figure 13-2 on page 13-3) by using the MAC authentication bypass feature. For example, you can enable this feature on 802.1x ports connected to devices such as printers.

If 802.1x authentication times out while waiting for an EAPOL response from the client, the switch tries to authorize the client by using MAC authentication bypass.

When the MAC authentication bypass feature is enabled on an 802.1x port, the switch uses the MAC address as the client identity. The authentication server has a database of client MAC addresses that are allowed network access. After detecting a client on an 802.1x port, the switch waits for an Ethernet packet from the client. The switch sends the authentication server a RADIUS-access/request frame with a username and password based on the MAC address. If authorization succeeds, the switch grants the client access to the network. If authorization fails, the switch assigns the port to the guest VLAN if one is configured.

If an EAPOL packet is detected on the interface during the lifetime of the link, the switch determines that the device connected to that interface is an 802.1x-capable supplicant and uses 802.1x authentication (not MAC authentication bypass) to authorize the interface. EAPOL history is cleared if the interface link status goes down.

If the switch already authorized a port by using MAC authentication bypass and detects an 802.1x supplicant, the switch does not unauthorize the client connected to the port. When reauthentication occurs, the switch uses 802.1x authentication as the preferred reauthentication process if the previous session ended because the Termination-Action RADIUS attribute value is DEFAULT.

Clients that were authorized with MAC authentication bypass can be reauthenticated. The reauthentication process is the same as that for clients that were authenticated with 802.1x. During reauthentication, the port remains in the previously assigned VLAN. If reauthentication is successful, the switch keeps the port in the same VLAN. If reauthentication fails, the switch assigns the port to the guest VLAN, if one is configured.

If reauthentication is based on the Session-Timeout RADIUS attribute (Attribute[27]) and the Termination-Action RADIUS attribute (Attribute [29]) and if the Termination-Action RADIUS attribute (Attribute [29]) action is *Initialize,* (the attribute value is *DEFAULT*), the MAC authentication bypass session ends, and connectivity is lost during reauthentication. If MAC authentication bypass is enabled and the 802.1x authentication times out, the switch uses the MAC authentication bypass feature to initiate reauthorization. For more information about these AV pairs, see RFC 3580, “802.1X Remote Authentication Dial In User Service (RADIUS) Usage Guidelines.”

MAC authentication bypass interacts with the features:

- 802.1x authentication—You can enable MAC authentication bypass only if 802.1x authentication is enabled on the port.
- Guest VLAN—If a client has an invalid MAC address identity, the switch assigns the client to a guest VLAN if one is configured.
- Restricted VLAN—This feature is not supported when the client connected to an 802.1x port is authenticated with MAC authentication bypass.
- Port security—See the “802.1x Authentication with Port Security” section on page 13-24.
- Voice VLAN—See the “802.1x Authentication with Voice VLAN Ports” section on page 13-23.
- VLAN Membership Policy Server (VMPS)—802.1x and VMPS are mutually exclusive.
- Private VLAN—You can assign a client to a private VLAN.

- Network admission control (NAC) Layer 2 IP validation—This feature takes effect after an 802.1x port is authenticated with MAC authentication bypass, including hosts in the exception list.
- Network Edge Access Topology (NEAT)—MAB and NEAT are mutually exclusive. You cannot enable MAB when NEAT is enabled on an interface, and you cannot enable NEAT when MAB is enabled on an interface.

For more configuration information, see the "Authentication Manager" section on page 13-6.

Cisco IOS Release 12.2(55)SE and later supports filtering of verbose MAB system messages. See the "Authentication Manager CLI Commands" section on page 13-8.

802.1x User Distribution

You can configure 802.1x user distribution to load-balance users with the same group name across multiple different VLANs.

The VLANs are either supplied by the RADIUS server or configured through the switch CLI under a VLAN group name.

- Configure the RADIUS server to send more than one VLAN name for a user. The multiple VLAN names can be sent as part of the response to the user. The 802.1x user distribution tracks all the users in a particular VLAN and achieves load balancing by moving the authorized user to the least populated VLAN.
- Configure the RADIUS server to send a VLAN group name for a user. The VLAN group name can be sent as part of the response to the user. You can search for the selected VLAN group name among the VLAN group names that you configured by using the switch CLI. If the VLAN group name is found, the corresponding VLANs under this VLAN group name are searched to find the least populated VLAN. Load balancing is achieved by moving the corresponding authorized user to that VLAN.

Note The RADIUS server can send the VLAN information in any combination of VLAN-IDs, VLAN names, or VLAN groups.

802.1x User Distribution Configuration Guidelines

- Confirm that at least one VLAN is mapped to the VLAN group.
- You can map more than one VLAN to a VLAN group.
- You can modify the VLAN group by adding or deleting a VLAN.
- When you clear an existing VLAN from the VLAN group name, none of the authenticated ports in the VLAN are cleared, but the mappings are removed from the existing VLAN group.
- If you clear the last VLAN from the VLAN group name, the VLAN group is cleared.
- You can clear a VLAN group even when the active VLANs are mapped to the group. When you clear a VLAN group, none of the ports or users that are in the authenticated state in any VLAN within the group are cleared, but the VLAN mappings to the VLAN group are cleared.

For more information, see the "Configuring 802.1x User Distribution" section on page 13-46.

Network Admission Control Layer 2 802.1x Validation

The switch supports the Network Admission Control (NAC) Layer 2 802.1x validation, which checks the antivirus condition or *posture* of endpoint systems or clients before granting the devices network access. With NAC Layer 2 802.1x validation, you can do these tasks:

- Download the Session-Timeout RADIUS attribute (Attribute[27]) and the Termination-Action RADIUS attribute (Attribute[29]) from the authentication server.
- Set the number of seconds between reauthentication attempts as the value of the Session-Timeout RADIUS attribute (Attribute[27]) and get an access policy against the client from the RADIUS server.
- Set the action to be taken when the switch tries to reauthenticate the client by using the Termination-Action RADIUS attribute (Attribute[29]). If the value is the *DEFAULT* or is not set, the session ends. If the value is RADIUS-Request, the reauthentication process starts.
- Set the list of VLAN number or name or VLAN group name as the value of the Tunnel Group Private ID (Attribute[81]) and the preference for the VLAN number or name or VLAN group name as the value of the Tunnel Preference (Attribute[83]). If you do not configure the Tunnel Preference, the first Tunnel Group Private ID (Attribute[81]) attribute is picked up from the list.
- View the NAC posture token, which shows the posture of the client, by using the **show authentication** privileged EXEC command.
- Configure secondary private VLANs as guest VLANs.

Configuring NAC Layer 2 802.1x validation is similar to configuring 802.1x port-based authentication except that you must configure a posture token on the RADIUS server. For information about configuring NAC Layer 2 802.1x validation, see the "Configuring NAC Layer 2 802.1x Validation" section on page 13-46 and the "Configuring Periodic Reauthentication" section on page 13-39.

For more information about NAC, see the *Network Admission Control Software Configuration Guide*.

For more configuration information, see the "Authentication Manager" section on page 13-6.

Flexible Authentication Ordering

You can use flexible authentication ordering to configure the order of methods that a port uses to authenticate a new host. MAC authentication bypass and 802.1x can be the primary or secondary authentication methods, and web authentication can be the fallback method if either or both of those authentication attempts fail. For the configuration commands, see "Configuring Optional 802.1x Authentication Features" section on page 13-40

Open1x Authentication

Open1x authentication allows a device access to a port before that device is authenticated. When open authentication is configured, a new host can pass traffic according to the access control list (ACL) defined on the port. After the host is authenticated, the policies configured on the RADIUS server are applied to that host.

You can configure open authentication with these scenarios:

- Single-host mode with open authentication—Only one user is allowed network access before and after authentication.
- MDA mode with open authentication—Only one user in the voice domain and one user in the data domain are allowed.
- Multiple-hosts mode with open authentication—Any host can access the network.
- Multiple-authentication mode with open authentication—Similar to MDA, except multiple hosts can be authenticated.

For more information see the "Configuring the Host Mode" section on page 13-38.

Note If open authentication is configured, it takes precedence over other authentication controls. This means that if you use the **authentication open** interface configuration command, the port will grant access to the host irrespective of the **authentication port-control** interface configuration command.

802.1x Supplicant and Authenticator Switches with Network Edge Access Topology (NEAT)

The Network Edge Access Topology (NEAT) feature extends identity to areas outside the wiring closet (such as conference rooms). This allows any type of device to authenticate on the port.

- You can configure a switch to act as a supplicant to another switch by using the 802.1x supplicant feature. This configuration is helpful in a scenario, where, for example, a switch is outside a wiring closet and is connected to an upstream switch through a trunk port. A switch configured with the 802.1x switch supplicant feature authenticates with the upstream switch for secure connectivity.

  Once the supplicant switch authenticates successfully the port mode changes from access to trunk.
- If the access VLAN is configured on the authenticator switch, it becomes the native VLAN for the trunk port after successful authentication.

You can enable MDA or multiauth mode on the authenticator switch interface that connects to one more supplicant switches. Multihost mode is not supported on the authenticator switch interface.

Use the **dot1x supplicant force-multicast** global configuration command on the supplicant switch for Network Edge Access Topology (NEAT) to work in all host modes.

- Host authorization ensures that only traffic from authorized hosts (connecting to the switch with supplicant) is allowed on the network. The switches use Client Information Signalling Protocol (CISP) to send the MAC addresses connecting to the supplicant switch to the authenticator switch, as shown in Figure 13-6.
- Auto enablement automatically enables trunk configuration on the authenticator switch, allowing user traffic from multiple VLANs coming from supplicant switches. Configure the cisco-av-pair as *device-traffic-class=switch* at the ACS. (You can configure this under the *group* or the *user* settings.)

Figure 13-6 Authenticator and Supplicant Switch using CISP

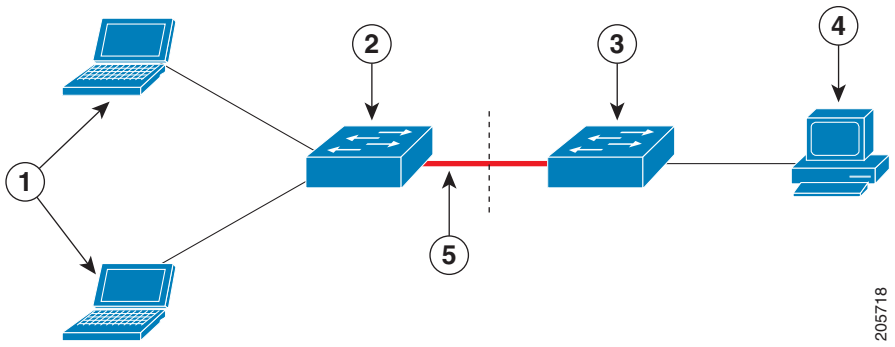

1	Workstations (clients)	**2**	Supplicant switch (outside wiring closet)
3	Authenticator switch	**4**	Access control server (ACS)
5	Trunk port		

802.1x Supplicant and Authenticator Switch Guidelines

- You can configure NEAT ports with the same configurations as the other authentication ports. When the supplicant switch authenticates, the port mode is changed from *access* to *trunk* based on the switch vendor-specific attributes (VSAs). (device-traffic-class=switch)
- The VSA changes the authenticator switch port mode from access to trunk and enables 802.1x trunk encapsulation and the access VLAN if any would be converted to a native trunk VLAN. VSA does not change any of the port configurations on the supplicant
- To change the host mode *and* to apply a standard port configuration on the authenticator switch port, you can also use Auto Smartports user-defined macros, instead of the switch VSA. This allows you to remove unsupported configurations on the authenticator switch port and to change the port mode from *access* to *trunk*. For information, see the *AutoSmartports Configuration Guide*.

For more information, see the "Configuring an Authenticator" section on page 13-47.

Using IEEE 802.1x Authentication with ACLs and the RADIUS Filter-Id Attribute

The switch supports both IP standard and IP extended port access control lists (ACLs) applied to ingress ports.

- ACLs that you configure
- ACLs from the Access Control Server (ACS)

An IEEE 802.1x port in single-host mode uses ACLs from the ACS to provide different levels of service to an IEEE 802.1x-authenticated user. When the RADIUS server authenticates this type of user and port, it sends ACL attributes based on the user identity to the switch. The switch applies the attributes to the port for the duration of the user session. If the session is over, authentication fails, or a link fails, the port becomes unauthorized, and the switch removes the ACL from the port.

Only IP standard and IP extended port ACLs from the ACS support the Filter-Id attribute. It specifies the name or number of an ACL. The Filter-id attribute can also specify the direction (inbound or outbound) and a user or a group to which the user belongs.

- The Filter-Id attribute for the user takes precedence over that for the group.
- If a Filter-Id attribute from the ACS specifies an ACL that is already configured, it takes precedence over a user-configured ACL.
- If the RADIUS server sends more than one Filter-Id attribute, only the last attribute is applied.

If the Filter-Id attribute is not defined on the switch, authentication fails, and the port returns to the unauthorized state.

Authentication Manager Common Session ID

Authentication manager uses a single session ID (referred to as a common session ID) for a client no matter which authentication method is used. This ID is used for all reporting purposes, such as the **show** commands and MIBs. The session ID appears with all per-session syslog messages.

The session ID includes:

- The IP address of the Network Access Device (NAD)
- A monotonically increasing unique 32-bit integer
- The session start time stamp (a 32-bit integer)

Default 802.1x Authentication Settings

Table 13-3 shows the default 802.1x authentication settings.

Table 13-3 Default 802.1x Authentication Settings

Feature	Default Setting
Switch 802.1x enable state	Disabled.
Per-port 802.1x enable state	Disabled (force-authorized). The port sends and receives normal traffic without 802.1x-based authentication of the client.
AAA	Disabled.
RADIUS server • IP address • UDP authentication port • Key	 • None specified. • 1812. • None specified.
Host mode	Single-host mode.
Control direction	Bidirectional control.

Table 13-3 Default 802.1x Authentication Settings (continued)

Feature	Default Setting
Periodic reauthentication	Disabled.
Number of seconds between reauthentication attempts	3600 seconds.
Reauthentication number	2 times (number of times that the switch restarts the authentication process before the port changes to the unauthorized state).
Quiet period	60 seconds (number of seconds that the switch remains in the quiet state following a failed authentication exchange with the client).
Retransmission time	30 seconds (number of seconds that the switch should wait for a response to an EAP request/identity frame from the client before resending the request).
Maximum retransmission number	2 times (number of times that the switch will send an EAP-request/identity frame before restarting the authentication process).
Client timeout period	30 seconds (when relaying a request from the authentication server to the client, the amount of time the switch waits for a response before resending the request to the client.)
Authentication server timeout period	30 seconds (when relaying a response from the client to the authentication server, the amount of time the switch waits for a reply before resending the response to the server.) You can change this timeout period by using the **authentication timer server** interface configuration command.
Inactivity timeout	Disabled.
Guest VLAN	None specified.
Inaccessible authentication bypass	Disabled.
Restricted VLAN	None specified.
Authenticator (switch) mode	None specified.
MAC authentication bypass	Disabled.
Voice-aware security	Disabled

802.1x Accounting

Enabling AAA system accounting with 802.1x accounting allows system reload events to be sent to the accounting RADIUS server for logging. The server can then infer that all active 802.1x sessions are closed.

Because RADIUS uses the unreliable UDP transport protocol, accounting messages might be lost due to poor network conditions. If the switch does not receive the accounting response message from the RADIUS server after a configurable number of retransmissions of an accounting request, this system message appears:

```
Accounting message %s for session %s failed to receive Accounting Response.
```

When the stop message is not sent successfully, this message appears:

```
00:09:55: %RADIUS-4-RADIUS_DEAD: RADIUS server 172.20.246.201:1645,1646 is not responding.
```

Note You must configure the RADIUS server to perform accounting tasks, such as logging start, stop, and interim-update messages and time stamps. To turn on these functions, enable logging of "Update/Watchdog packets from this AAA client" in your RADIUS server Network Configuration tab. Next, enable "CVS RADIUS Accounting" in your RADIUS server System Configuration tab.

802.1x Authentication Guidelines

- When 802.1x authentication is enabled, ports are authenticated before any other Layer 2 features are enabled.
- If the VLAN to which an 802.1x-enabled port is assigned changes, this change is transparent and does not affect the switch. For example, this change occurs if a port is assigned to a RADIUS server-assigned VLAN and is then assigned to a different VLAN after reauthentication.

  If the VLAN to which an 802.1x port is assigned to shut down, disabled, or removed, the port becomes unauthorized. For example, the port is unauthorized after the access VLAN to which a port is assigned shuts down or is removed.
- The 802.1x protocol is supported on Layer 2 static-access ports, and voice VLAN ports, but it is not supported on these port types:
  - Trunk port—If you try to enable 802.1x authentication on a trunk port, an error message appears, and 802.1x authentication is not enabled. If you try to change the mode of an 802.1x-enabled port to trunk, an error message appears, and the port mode is not changed.
  - Dynamic ports—A port in dynamic mode can negotiate with its neighbor to become a trunk port. If you try to enable 802.1x authentication on a dynamic port, an error message appears, and 802.1x authentication is not enabled. If you try to change the mode of an 802.1x-enabled port to dynamic, an error message appears, and the port mode is not changed.
  - Dynamic-access ports—If you try to enable 802.1x authentication on a dynamic-access (VLAN Query Protocol [VQP]) port, an error message appears, and 802.1x authentication is not enabled. If you try to change an 802.1x-enabled port to dynamic VLAN assignment, an error message appears, and the VLAN configuration is not changed.
  - EtherChannel port—Do not configure a port that is an active or a not-yet-active member of an EtherChannel as an 802.1x port. If you try to enable 802.1x authentication on an EtherChannel port, an error message appears, and 802.1x authentication is not enabled.
  - Switched Port Analyzer (SPAN) and Remote SPAN (RSPAN) destination ports—You can enable 802.1x authentication on a port that is a SPAN or RSPAN destination port. However, 802.1x authentication is disabled until the port is removed as a SPAN or RSPAN destination port. You can enable 802.1x authentication on a SPAN or RSPAN source port.
- Before globally enabling 802.1x authentication on a switch by entering the **dot1x system-auth-control** global configuration command, remove the EtherChannel configuration from the interfaces on which 802.1x authentication and EtherChannel are configured.
- System messages related to 802.1x authentication can be filtered. See the "Authentication Manager CLI Commands" section on page 13-8.

VLAN Assignment, Guest VLAN, Restricted VLAN, and Inaccessible Authentication Bypass Guidelines

- When 802.1x authentication is enabled on a port, you cannot configure a port VLAN that is equal to a voice VLAN.
- The 802.1x authentication with VLAN assignment feature is not supported on trunk ports, dynamic ports, or with dynamic-access port assignment through a VMPS.
- You can configure 802.1x authentication on a private-VLAN port, but do not configure 802.1x authentication with port security, a voice VLAN, a guest VLAN, a restricted VLAN, or a per-user ACL on private-VLAN ports.
- You can configure any VLAN except an RSPAN VLAN, private VLAN, or a voice VLAN as an 802.1x guest VLAN. The guest VLAN feature is not supported on internal VLANs (routed ports) or trunk ports; it is supported only on access ports.
- After you configure a guest VLAN for an 802.1x port to which a DHCP client is connected, you might need to get a host IP address from a DHCP server. You can change the settings for restarting the 802.1x authentication process on the switch before the DHCP process on the client times out and tries to get a host IP address from the DHCP server. Decrease the settings for the 802.1x authentication process (**authentication timer inactivity** and **authentication timer reauthentication** interface configuration commands). The amount to decrease the settings depends on the connected 802.1x client type.
- When configuring the inaccessible authentication bypass feature, follow these guidelines:
  - The feature is supported on 802.1x port in single-host mode and multihosts mode.
  - If the client is running Windows XP and the port to which the client is connected is in the critical-authentication state, Windows XP might report that the interface is not authenticated.
  - If the Windows XP client is configured for DHCP and has an IP address from the DHCP server, receiving an EAP-Success message on a critical port might not reinitiate the DHCP configuration process.
  - You can configure the inaccessible authentication bypass feature and the restricted VLAN on an 802.1x port. If the switch tries to reauthenticate a critical port in a restricted VLAN and all the RADIUS servers are unavailable, switch changes the port state to the critical authentication state and remains in the restricted VLAN.
- You can configure any VLAN except an RSPAN VLAN or a voice VLAN as an 802.1x restricted VLAN. The restricted VLAN feature is not supported on internal VLANs (routed ports) or trunk ports; it is supported only on access ports.

MAC Authentication Bypass Guidelines

- Unless otherwise stated, the MAC authentication bypass guidelines are the same as the 802.1x authentication guidelines. For more information, see the "802.1x Authentication Guidelines" section on page 13-32.
- If you disable MAC authentication bypass from a port after the port has been authorized with its MAC address, the port state is not affected.
- If the port is in the unauthorized state and the client MAC address is not the authentication-server database, the port remains in the unauthorized state. However, if the client MAC address is added to the database, the switch can use MAC authentication bypass to reauthorize the port.

- If the port is in the authorized state, the port remains in this state until reauthorization occurs.
- You can configure a timeout period for hosts that are connected by MAC authentication bypass but are inactive. The range is 1to 65535 seconds.

Maximum Number of Allowed Devices Per Port Guidelines

This is the maximum number of devices allowed on an 802.1x-enabled port:

- In single-host mode, only one device is allowed on the access VLAN. If the port is also configured with a voice VLAN, an unlimited number of Cisco IP phones can send and receive traffic through the voice VLAN.
- In multidomain authentication (MDA) mode, one device is allowed for the access VLAN, and one IP phone is allowed for the voice VLAN.
- In multiple-host mode, only one 802.1x supplicant is allowed on the port, but an unlimited number of non-802.1x hosts are allowed on the access VLAN. An unlimited number of devices are allowed on the voice VLAN.

How to Configure IEEE 802.1x Port-Based Authentication

802.1x Authentication Configuration Process

To configure 802.1x port-based authentication, you must enable authentication, authorization, and accounting (AAA) and specify the authentication method list. A method list describes the sequence and authentication method to be queried to authenticate a user.

To allow per-user ACLs or VLAN assignment, you must enable AAA authorization to configure the switch for all network-related service requests.

This is the 802.1x AAA configuration process:

Step 1 A user connects to a port on the switch.

Step 2 Authentication is performed.

Step 3 The VLAN assignment is enabled, as appropriate, based on the RADIUS server configuration.

Step 4 The switch sends a start message to an accounting server.

Step 5 Reauthentication is performed, as necessary.

Step 6 The switch sends an interim accounting update to the accounting server, that is based on the result of reauthentication.

Step 7 The user disconnects from the port.

Step 8 The switch sends a stop message to the accounting server.

Beginning in privileged EXEC mode, follow these steps to configure 802.1x port-based authentication:

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**aaa new-model**	Enables AAA.
Step 3	**aaa authentication dot1x** {**default**} *method1*	Creates an 802.1x authentication method list. To create a default list to use when a named list is *not* specified in the **authentication** command, use the **default** keyword followed by the method to use in default situations. The default method list is automatically applied to all ports. For *method1*, enter the **group radius** keywords to use the list of all RADIUS servers for authentication. **Note** Though other keywords are visible in the command-line help string, only the **group radius** keywords are supported.
Step 4	**dot1x system-auth-control**	Enables 802.1x authentication globally on the switch.
Step 5	**aaa authorization network** {**default**} **group radius**	(Optional) Configures the switch to use user-RADIUS authorization for all network-related service requests, such as per-user ACLs or VLAN assignment. For per-user ACLs, single-host mode must be configured. This setting is the default.
Step 6	**radius-server host** *ip-address*	(Optional) Specifies the IP address of the RADIUS server.
Step 7	**radius-server key** *string*	(Optional) Specifies the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server.
Step 8	**interface** *interface-id*	Specifies the port connected to the client to enable for 802.1x authentication, and enter interface configuration mode.
Step 9	**switchport mode access**	(Optional) Sets the port to access mode only if you configured the RADIUS server in Step 6 and Step 7.
Step 10	**authentication port-control auto**	Enables 802.1x authentication on the port.
Step 11	**dot1x pae authenticator**	Sets the interface Port Access Entity to act only as an authenticator and ignore messages meant for a supplicant.
Step 12	**end**	Returns to privileged EXEC mode.
Step 13	**show authentication**	Verifies your entries.
Step 14	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring the Switch-to-RADIUS-Server Communication

You can globally configure the timeout, retransmission, and encryption key values for all RADIUS servers by using the **radius-server host** global configuration command. If you want to configure these options on a per-server basis, use the **radius-server timeout**, **radius-server retransmit**, and the **radius-server key** global configuration commands. For more information, see the "Configuring Settings for All RADIUS Servers" section on page 12-37.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**radius-server host** {*hostname* | *ip-address*} **auth-port** *port-number* **key** *string*	Configures the RADIUS server parameters. *hostname* | *ip-address*—Specifies the hostname or IP address of the remote RADIUS server. **auth-port** *port-number*—Specifies the UDP destination port for authentication requests. The default is 1812. The range is 0 to 65536. **key** *string*—Specifies the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. The key is a text string that must match the encryption key used on the RADIUS server. **Note** Always configure the key as the last item in the **radius-server host** command syntax because leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in the key, do not enclose the key in quotation marks unless the quotation marks are part of the key. This key must match the encryption used on the RADIUS daemon. If you want to use multiple RADIUS servers, reenter this command.
Step 3	**end**	Returns to privileged EXEC mode.
Step 4	**show running-config**	Verifies your entries.
Step 5	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring 802.1x Readiness Check

	Command	Purpose
Step 1	**dot1x test eapol-capable** [**interface** *interface-id*]	Enables the 802.1x readiness check on the switch. *interface-id*—Specifies the port on which to check for 802.1x readiness. **Note** If you omit the optional **interface** keyword, all interfaces on the switch are tested.
Step 1	**configure terminal**	(Optional) Enters global configuration mode.
Step 2	**dot1x test timeout** *timeout*	(Optional) Configures the timeout used to wait for EAPOL response. The range is from 1 to 65535 seconds. The default is 10 seconds.
Step 3	**end**	(Optional) Returns to privileged EXEC mode.
Step 4	**show running-config**	(Optional) Verifies your modified timeout values.

Enabling Voice Aware 802.1x Security

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**errdisable detect cause security-violation shutdown vlan**	Shuts down any VLAN on which a security violation error occurs. **Note** If the **shutdown vlan** keywords are not included, the entire port enters the error-disabled state and shuts down.
Step 3	**errdisable recovery cause security-violation**	(Optional) Enables automatic per-VLAN error recovery.
Step 4	**clear errdisable interface** *interface-id* **vlan** [*vlan-list*]	(Optional) Reenables individual VLANs that have been error-disabled. • *interface-id—S*pecifies the port on which to reenable individual VLANs. • (Optional) *vlan-list—S*pecifies a list of VLANs to be reenabled. If *vlan-list* is not specified, all VLANs are reenabled.
Step 5	**shutdown** **no-shutdown**	(Optional) Reenables an error-disabled VLAN, and clear all error-disable indications.
Step 6	**end**	Returns to privileged EXEC mode.
Step 7	**show errdisable detect**	Verifies your entries.
Step 8	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring 802.1x Violation Modes

You can configure an 802.1x port so that it shuts down, generates a syslog error, or discards packets from a new device when:

- A device connects to an 802.1x-enabled port
- The maximum number of allowed about devices have been authenticated on the port

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**aaa new-model**	Enables AAA.
Step 3	**aaa authentication dot1x** {**default**} *method1*	Creates an 802.1x authentication method list. To create a default list to use when a named list is *not* specified in the **authentication** command, use the **default** keyword followed by the method that is to be used in default situations. The default method list is automatically applied to all ports. *method1—S*pecifies the **group radius** keywords to use the list of all RADIUS servers for authentication. **Note** Though other keywords are visible in the command-line help string, only the **group radius** keywords are supported.
Step 4	**interface** *interface-id*	Specifies the port connected to the client that is to be enabled for 802.1x authentication, and enter interface configuration mode.

	Command	Purpose
Step 5	**switchport mode access**	Sets the port to access mode.
Step 6	**authentication violation** {**shutdown** | **restrict** | **protect** | **replace**}	Configures the violation mode. • **shutdown**—Error-disables the port. • **restrict**—Generates a syslog error. • **protect**—Drops packets from any new device that sends traffic to the port. • **replace**—Removes the current session and authenticates with the new host.
Step 7	**end**	Returns to privileged EXEC mode.
Step 8	**show authentication**	Verifies your entries.
Step 9	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring the Host Mode

This task describes how to configure a single host (client) or multiple hosts on an 802.1x-authorized port.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**radius-server vsa send authentication**	Configures the network access server to recognize and use vendor-specific attributes (VSAs).
Step 3	**interface** *interface-id*	Specifies the port to which multiple hosts are indirectly attached, and enter interface configuration mode.
Step 4	**authentication host-mode** [**multi-auth** | **multi-domain** | **multi-host** | **single-host**]	The keywords have these meanings: • **multi-auth**—Allows one client on the voice VLAN and multiple authenticated clients on the data VLAN. Each host is individually authenticated. **Note** The **multi-auth** keyword is only available with the **authentication host-mode** command. • **multi-host**—Allows multiple hosts on an 802.1x-authorized port after a single host has been authenticated. • **multi-domain**—Allows both a host and a voice device, such as an IP phone (Cisco or non-Cisco), to be authenticated on an 802.1x-authorized port. **Note** You must configure the voice VLAN for the IP phone when the host mode is set to **multi-domain**. For more information, see Chapter 19, "Configuring Voice VLAN." • **single-host**—Allows a single host (client) on an 802.1x-authorized port. Make sure that the **authentication port-control** interface configuration command set is set to **auto** for the specified interface.

	Command	Purpose
Step 5	**switchport voice vlan** *vlan-id*	(Optional) Configures the voice VLAN.
Step 6	**end**	Returns to privileged EXEC mode.
Step 7	**show authentication interface** *interface-id*	Verifies your entries.
Step 8	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring Periodic Reauthentication

You can enable periodic 802.1x client reauthentication and specify how often it occurs. If you do not specify a time period before enabling reauthentication, the number of seconds between attempts is 3600. Beginning in privileged EXEC mode, follow these steps to enable periodic reauthentication of the client and to configure the number of seconds between reauthentication attempts. This procedure is optional.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the port to be configured, and enter interface configuration mode.
Step 3	**authentication periodic**	Enables periodic reauthentication of the client, which is disabled by default. **Note** The default value is 3600 seconds. To change the value of the reauthentication timer or to have the switch use a RADIUS-provided session timeout, enter the **authentication timer reauthenticate** command.
Step 4	**authentication timer** {{[**inactivity** | **reauthenticate**]} {**restart** *value*}}	Sets the number of seconds between reauthentication attempts. • **inactivity**—Interval in seconds after which if there is no activity from the client then it is unauthorized • **reauthenticate**—Time in seconds after which an automatic reauthentication attempt is be initiated. • **restart** *value*—Interval in seconds after which an attempt is made to authenticate an unauthorized port. This command affects the behavior of the switch only if periodic reauthentication is enabled.
Step 5	**authentication timer reauthenticate** *seconds*	Sets the number of seconds that the switch waits for a response to an EAP-request/identity frame from the client before resending the request. The range is 1 to 65535 seconds; the default is 5. **Note** You should change the default value of this command only to adjust for unusual circumstances such as unreliable links or specific behavioral problems with certain clients and authentication servers.
Step 6	**end**	Returns to privileged EXEC mode.

	Command	Purpose
Step 7	**show authentication interface** *interface-id*	Verifies your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Configuring Optional 802.1x Authentication Features

	Command	Purpose
Step 1	**dot1x reauthenticate interface** ***interface-id***	(Optional) Manually initiates a reauthentication of the specified IEEE 802.1x-enabled port.
Step 2	**authentication mac-move permit**	(Optional) Enables MAC move on the switch.
Step 3	**authentication violation** {**protect** | **replace** | **restrict** | **shutdown**}	(Optional) **replace**—Enables MAC replace on the interface. The port removes the current session and initiates authentication with the new host. The other keywords have these effects: • **protect**—Drops port packets with unexpected MAC addresses without generating a system message. • **restrict**—Drops violating packets by the CPU and a system message is generated. • **shutdown**—Error-disables the port when it receives an unexpected MAC address.
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mab request format attribute 32 vlan access-vlan**	(Optional) Enables VLAN ID-based MAC authentication.
Step 3	**interface** *interface-id*	(Optional) Specifies the port to be configured, and enters interface configuration mode.
Step 4	**authentication timer inactivity** *seconds*	(Optional) Sets the number of seconds that the switch remains in the quiet state after a failed authentication exchange with the client. The range is 1 to 65535 seconds; the default is 60.
Step 5	**authentication timer reauthenticate** *seconds*	(Optional) Sets the number of seconds that the switch waits for a response to an EAP-request/identity frame from the client before resending the request. The range is 1 to 65535 seconds; the default is 5. **Note** You should change the default value of this command only to adjust for unusual circumstances such as unreliable links or specific behavioral problems with certain clients and authentication servers.

	Command	Purpose
Step 6	**dot1x max-reauth-req** *count*	(Optional) Sets the number of times that the switch sends an EAP-request/identity frame to the client before restarting the authentication process. The range is 1 to 10; the default is 2. **Note** You should change the default value of this command only to adjust for unusual circumstances such as unreliable links or specific behavioral problems with certain clients and authentication servers.
Step 7	**dot1x max-req** *count*	(Optional) Sets the number of times that the switch restarts the authentication process before the port changes to the unauthorized state. The range is 0 to 10; the default is 2.
Step 8	**authentication control-direction {both | in}**	(Optional) Enables 802.1x authentication with WoL on the port, and uses these keywords to configure the port as bidirectional or unidirectional. • **both**—Sets the port as bidirectional. The port cannot receive packets from or send packets to the host. By default, the port is bidirectional. • **in**—Sets the port as unidirectional. The port can send packets to the host but cannot receive packets from the host.
Step 9	**authentication order** [**mab**] {**webauth**}	(Optional) Sets the order of authentication methods. • **mab**—Adds MAC authentication bypass (MAB) to the order of authentication methods. • **webauth**—Adds web authentication to the order of authentication methods.
Step 10	**authentication order** [**dot1x** | **mab**] | {**webauth**}	(Optional) Sets the order of authentication methods used on a port.
Step 11	**authentication priority** [**dot1x** | **mab**] | {**webauth**}	(Optional) Adds an authentication method to the port-priority list.
Step 12	**dot1x default**	Resets the 802.1x parameters to the default values.
Step 13	**end**	Returns to privileged EXEC mode.
Step 14	**show authentication interface** *interface-id*	Verifies your entries.
Step 15	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring 802.1x Accounting

Before You Begin

AAA must be enabled on your switch.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the port to be configured, and enter interface configuration mode.
Step 3	**aaa accounting dot1x default start-stop group radius**	Enables 802.1x accounting using the list of all RADIUS servers.
Step 4	**aaa accounting system default start-stop group radius**	(Optional) Enables system accounting (using the list of all RADIUS servers) and generates system accounting reload event messages when the switch reloads.
Step 5	**end**	Returns to privileged EXEc mode.
Step 6	**show running-config**	Verifies your entries.
Step 7	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring a Guest VLAN

When you configure a guest VLAN, clients that are not 802.1x-capable are put into the guest VLAN when the server does not receive a response to its EAP request/identity frame. Clients that are 802.1x-capable but that fail authentication are not granted network access. The switch supports guest VLANs in single-host or multiple-hosts mode.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the port to be configured, and enters interface configuration mode.
Step 3	**switchport mode access**	Sets the port to access mode
	or	or
	switchport mode private-vlan host	Configures the Layer 2 port as a private-VLAN host port.
Step 4	**authentication port-control auto**	Enables 802.1x authentication on the port.
Step 5	**authentication event no-response action authorize vlan** *vlan-id*	Specifies an active VLAN as an 802.1x guest VLAN. The range is 1 to 4096. You can configure any active VLAN except an internal VLAN (routed port), an RSPAN VLAN, a primary private VLAN, or a voice VLAN as an 802.1x guest VLAN.
Step 6	**end**	Returns to privileged EXEC mode.
Step 7	**show authentication interface** *interface-id*	Verifies your entries.
Step 8	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring a Restricted VLAN

When you configure a restricted VLAN on a switch, clients that are 802.1x-compliant are moved into the restricted VLAN when the authentication server does not receive a valid username and password. The switch supports restricted VLANs only in single-host mode.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the port to be configured, and enters interface configuration mode.
Step 3	**switchport mode access**	Sets the port to access mode,
	or	or
	switchport mode private-vlan host	Configures the Layer 2 port as a private-VLAN host port.
Step 4	**authentication port-control auto**	Enables 802.1x authentication on the port.
Step 5	**authentication event fail action authorize** *vlan-id*	Specifies an active VLAN as an 802.1x restricted VLAN. The range is 1 to 4096. You can configure any active VLAN except an internal VLAN (routed port), an RSPAN VLAN, a primary private VLAN, or a voice VLAN as an 802.1x restricted VLAN.
Step 6	**end**	Returns to privileged EXEC mode.
Step 7	**show authentication interface** *interface-id*	(Optional) Verifies your entries.
Step 8	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring the Maximum Number of Authentication Attempts

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the port to be configured, and enters interface configuration mode.
Step 3	**switchport mode access**	Sets the port to access mode,
	or	or
	switchport mode private-vlan host	Configures the Layer 2 port as a private-VLAN host port.
Step 4	**authentication port-control auto**	Enables 802.1x authentication on the port.
Step 5	**authentication event fail action authorize** *vlan-id*	Specifies an active VLAN as an 802.1x restricted VLAN. The range is 1 to 4096. You can configure any active VLAN except an internal VLAN (routed port), an RSPAN VLAN, a primary private VLAN, or a voice VLAN as an 802.1x restricted VLAN.
Step 6	**authentication event retry** *retry count*	Specifies a number of authentication attempts to allow before a port moves to the restricted VLAN. The range is 1 to 3, and the default is 3.

	Command	Purpose
Step 7	**end**	Returns to privileged EXEC mode.
Step 8	**show authentication interface** *interface-id*	(Optional) Verifies your entries.
Step 9	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring Inaccessible Authentication Bypass

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**radius-server dead-criteria time** *time* **tries** *tries*	(Optional) Sets the conditions that are used to decide when a RADIUS server is considered unavailable or *dead*. The range for *time* is from 1 to 120 seconds. The switch dynamically determines the default *seconds* value that is 10 to 60 seconds. The range for *tries* is from 1 to 100. The switch dynamically determines the default *tries* parameter that is 10 to 100.
Step 3	**radius-server deadtime** *minutes*	(Optional) Sets the number of minutes that a RADIUS server is not sent requests. The range is from 0 to 1440 minutes (24 hours). The default is 0 minutes.

	Command	Purpose
Step 4	**radius-server host** *ip-address* **[acct-port** *udp-port*] [**auth-port** *udp-port*] [**test username** *name* [**idle-time** *time*] [**ignore-acct-port**] [**ignore-auth-port**]] [**key** *string*]	(Optional) Configures the RADIUS server parameters by using these keywords: • **acct-port** *udp-port*—Specifies the UDP port for the RADIUS accounting server. The range for the UDP port number is from 0 to 65536. The default is 1646. • **auth-port** *udp-port*—Specifies the UDP port for the RADIUS authentication server. The range for the UDP port number is from 0 to 65536. The default is 1645. **Note** You should configure the UDP port for the RADIUS accounting server and the UDP port for the RADIUS authentication server to nondefault values. • **test username** *name*—Enables automated testing of the RADIUS server status, and specifies the username to be used. • **idle-time** *time*—Sets the interval of time in minutes after which the switch sends test packets to the server. The range is from 1 to 35791 minutes. The default is 60 minutes (1 hour). • **ignore-acct-port**—Disables testing on the RADIUS-server accounting port. • **ignore-auth-port**—Disables testing on the RADIUS-server authentication port. • **key** *string*—Specifies the authentication and encryption key for all RADIUS communication between the switch and the RADIUS daemon. **Note** Always configure the key as the last item in the **radius-server host** command syntax because leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in the key, do not enclose the key in quotation marks unless the quotation marks are part of the key. This key must match the encryption used on the RADIUS daemon. You can also configure the authentication and encryption key by using the **radius-server key** {**0** *string* | **7** *string* | *string*} global configuration command.
Step 5	**dot1x critical** {**eapol** | **recovery delay** *milliseconds*}	(Optional) Configures the parameters for inaccessible authentication bypass. • **eapol**—Specifies that the switch sends an EAPOL-Success message when the switch successfully authenticates the critical port. • **recovery delay** *milliseconds*—Sets the recovery delay period during which the switch waits to reinitialize a critical port when a RADIUS server that was unavailable becomes available. The range is from 1 to 10000 milliseconds. The default is 1000 milliseconds (a port can be reinitialized every second).
Step 6	**interface** *interface-id*	Specifies the port to be configured, and enter interface configuration mode.
Step 7	**authentication event server dead action [authorize | reinitialize] vlan** *vlan-id*	Use these keywords to move hosts on the port if the RADIUS server is unreachable: • **authorize**—Moves any new hosts trying to authenticate to the user-specified critical VLAN. • **reinitialize**—Moves all authorized hosts on the port to the user-specified critical VLAN.
Step 8	**authentication event server dead action** {**authorize** | **reinitialize**} **vlan** *vlan-id*]	Enables the inaccessible authentication bypass feature and uses these keywords to configure the feature: • **authorize**—Authorizes the port. • **reinitialize**—Reinitializes all authorized clients.

	Command	Purpose
Step 9	**authentication server dead action authorize** [**vlan**]	Authorizes the switch in access VLAN or configured VLAN (if the VLAN is specified) when the ACS server is down.
Step 10	**end**	Returns to privileged EXEC mode.
Step 11	**show authentication interface** *interface-id*	(Optional) Verifies your entries.
Step 12	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring 802.1x User Distribution

Beginning in global configuration, follow these steps to configure a VLAN group and to map a VLAN to it:

	Command	Purpose
Step 1	**vlan group** *vlan-group-name* **vlan-list** *vlan-list*	Configures a VLAN group, and maps a single VLAN or a range of VLANs to it.
Step 2	**show vlan group all** *vlan-group-name*	Verifies the configuration.
Step 3	**no vlan group** *vlan-group-name* **vlan-list** *vlan-list*	Clears the VLAN group configuration or elements of the VLAN group configuration.

Configuring NAC Layer 2 802.1x Validation

You can configure NAC Layer 2 802.1x validation, which is also referred to as 802.1x authentication with a RADIUS server.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the port to be configured, and enters interface configuration mode.
Step 3	**authentication event no-response action authorize vlan** *vlan-id*	Specifies an active VLAN as an 802.1x guest VLAN. The range is 1 to 4096. You can configure any active VLAN except an internal VLAN (routed port), an RSPAN VLAN, or a voice VLAN as an 802.1x guest VLAN.
Step 4	**authentication periodic**	Enables periodic reauthentication of the client, which is disabled by default.
Step 5	**authentication timer reauthenticate**	Sets reauthentication attempt for the client (set to one hour). This command affects the behavior of the switch only if periodic reauthentication is enabled.
Step 6	**end**	Returns to privileged EXEC mode.

	Command	Purpose
Step 7	**show authentication interface** *interface-id*	Verifies your 802.1x authentication configuration.
Step 8	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring an Authenticator and Supplicant

You can also use an Auto Smartports user-defined macro instead of the switch VSA to configure the authenticator switch. For information, see the"Configuring Smartports Macros" chapter.

Configuring an Authenticator

Before You Begin

One switch outside a wiring closet must be configured as a supplicant and be connected to an authenticator switch.

Note The *cisco-av-pairs* must be configured as *device-traffic-class=switch* on the ACS, which sets the interface as a trunk after the supplicant is successfully authenticated.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**cisp enable**	Enables CISP.
Step 3	**interface** *interface-id*	Specifies the port to be configured, and enters interface configuration mode.
Step 4	**switchport mode access**	Sets the port mode to **access**.
Step 5	**authentication port-control auto**	Sets the port-authentication mode to auto.
Step 6	**dot1x pae authenticator**	Configures the interface as a port access entity (PAE) authenticator.
Step 7	**spanning-tree portfast**	Enables Port Fast on an access port connected to a single workstation or server.
Step 8	**end**	Returns to privileged EXEC mode.
Step 9	**show running-config interface** *interface-id*	Verifies your configuration.
Step 10	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring a Supplicant Switch with NEAT

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**cisp enable**	Enables CISP.

	Command	Purpose
Step 3	**dot1x credentials** *profile*	Creates 802.1x credentials profile. This must be attached to the port that is configured as supplicant.
Step 4	**username** *suppswitch*	Creates a username.
Step 5	**password** *password*	Creates a password for the new username.
Step 6	**dot1x supplicant force-multicast**	Forces the switch to send *only* multicast EAPOL packets when it receives either unicast or multicast packets. This also allows NEAT to work on the supplicant switch in all host modes.
Step 7	**interface** *interface-id*	Specifies the port to be configured, and enters interface configuration mode.
Step 8	**switchport trunk encapsulation dot1q**	Sets the port to trunk mode.
Step 9	**switchport mode trunk**	Configures the interface as a VLAN trunk port.
Step 10	**dot1x pae supplicant**	Configures the interface as a port access entity (PAE) supplicant.
Step 11	**dot1x credentials** *profile-name*	Attaches the 802.1x credentials profile to the interface.
Step 12	**end**	Returns to privileged EXEC mode.
Step 13	**show running-config interface** *interface-id*	Verifies your configuration.
Step 14	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring 802.1x Authentication with Downloadable ACLs and Redirect URLs

In addition to configuring 802.1x authentication on the switch, you need to configure the ACS. For more information, see the Cisco Secure ACS configuration guides.

Note You must configure a downloadable ACL on the ACS before downloading it to the switch.

Configuring Downloadable ACLs

The policies take effect after client authentication and the client IP address addition to the IP device tracking table. The switch then applies the downloadable ACL to the port.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip device tracking**	Configures the IP device tracking table.
Step 3	**aaa new-model**	Enables AAA.
Step 4	**aaa authorization network default group radius**	Sets the authorization method to local. To remove the authorization method, use the **no aaa authorization network default group radius** command.
Step 5	**radius-server vsa send authentication**	Configures the RADIUS VSA send authentication.

	Command	Purpose
Step 6	**interface** *interface-id*	Specifies the port to be configured, and enters interface configuration mode.
Step 7	**ip access-group** *acl-id* **in**	Configures the default ACL on the port in the input direction. **Note** The *acl-id* is an access list name or number.
Step 8	**show running-config interface** *interface-id*	Verifies your configuration.
Step 9	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring a Downloadable Policy

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**access-list** *access-list-number* **deny** *source* [*source-wildcard* **log**]	Defines the default port ACL by using a source address and wildcard. The access-list-number is a decimal number from 1 to 99 or 1300 to 1999. **deny** or **permit**—Specifies whether to deny or permit access if conditions are matched. *source*—Specifies the source address of the network or host that sends a packet: • The 32-bit quantity in dotted-decimal format. • The keyword **any** as an abbreviation for **source** and *source-wildcard* value of 0.0.0.0 255.255.255.255. You do not need to enter a *source-wildcard* value. • The keyword host as an abbreviation for **source** and *source-wildcard* of source 0.0.0.0. (Optional) *source-wildcard*—Applies the wildcard bits to the source. (Optional) **log**—Creates an informational logging message about the packet that matches the entry to be sent to the console.
Step 3	**interface** *interface-id*	Enters interface configuration mode.
Step 4	**ip access-group** *acl-id* **in**	Configures the default ACL on the port in the input direction. **Note** The *acl-id* is an access list name or number.
Step 5	**exit**	Returns to global configuration mode.
Step 6	**aaa new-model**	Enables AAA.
Step 7	**aaa authorization network default group radius**	Sets the authorization method to local. To remove the authorization method, use the **no aaa authorization network default group radius** command.
Step 8	**ip device tracking**	Enables the IP device tracking table. To disable the IP device tracking table, use the **no ip device tracking** global configuration commands.

	Command	Purpose
Step 9	**ip device tracking probe** [**count** | **interval** | **use-svi**]	(Optional) Configures the IP device tracking table: • **count** *count*—Sets the number of times that the switch sends the ARP probe. The range is from 1 to 5. The default is 3. • **interval** *interval*—Sets the number of seconds that the switch waits for a response before resending the ARP probe. The range is from 30 to 300 seconds. The default is 30 seconds. • **use-sv**i—Uses the switch virtual interface (SVI) IP address as source of ARP probes.
Step 10	**radius-server vsa send authentication**	Configures the network access server to recognize and uses vendor-specific attributes. **Note** The downloadable ACL must be operational.
Step 11	**end**	Returns to privileged EXEC mode.
Step 12	**show ip device tracking all**	Displays information about the entries in the IP device tracking table.
Step 13	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring Open1x

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the port to be configured, and enters interface configuration mode.
Step 3	**authentication control-direction {both | in}**	(Optional) Configures the port control as unidirectional or bidirectional.
Step 4	**authentication fallback** *name*	(Optional) Configures a port to use web authentication as a fallback method for clients that do not support 802.1x authentication.
Step 5	**authentication host-mode** [**multi-auth** | **multi-domain** | **multi-host** | **single-host**]	(Optional) Sets the authorization manager mode on a port.
Step 6	**authentication open**	(Optional) Enables or disables open access on a port.
Step 7	**authentication order** [**dot1x** | **mab**] | {**webauth**}	(Optional) Sets the order of authentication methods used on a port.
Step 8	**authentication periodic**	(Optional) Enables or disables reauthentication on a port.
Step 9	**authentication port-control** {**auto** | **force-authorized** | **force-un authorized**}	(Optional) Enables manual control of the port authorization state.
Step 10	**show authentication**	(Optional) Verifies your entries.
Step 11	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Resetting the 802.1x Authentication Configuration to the Default Values

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Enters interface configuration mode, and specifies the port to be configured.
Step 3	**dot1x default**	Resets the 802.1x parameters to the default values.
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show authentication interface** *interface-id*	Verifies your entries.
Step 6	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Monitoring and Maintaining IEEE 802.1x Port-Based Authentication

Command	Purpose
show dot1x all statistics	Displays 802.1x statistics for all ports.
show dot1x statistics interface *interface-id*	Displays 802.1x statistics for a specific port.
show dot1x all [**details** | **statistics** | **summary**]	Displays the 802.1x administrative and operational status for the switch.
show dot1x interface *interface-id*	Displays the 802.1x administrative and operational status for a specific port.

Configuration Examples for Configuring IEEE 802.1x Port-Based Authentication

Enabling a Readiness Check: Example

This example shows how to enable a readiness check on a switch to query a port. It also shows the response received from the queried port verifying that the device connected to it is 802.1x-capable:

```
switch# dot1x test eapol-capable interface gigabitethernet1/2

DOT1X_PORT_EAPOL_CAPABLE:DOT1X: MAC 00-01-02-4b-f1-a3 on gigabitethernet1/2 is EAPOL
capable
```

Enabling 802.1x Authentication: Example

This example shows how to enable 802.1x authentication and to allow multiple hosts:

```
Switch(config)# interface gigabitethernet1/2
Switch(config-if)# authentication port-control auto
Switch(config-if)# authentication host-mode multi-host
Switch(config-if)# end
```

Enabling MDA: Example

This example shows how to enable MDA and to allow both a host and a voice device on the port:

```
Switch(config)# interface gigabitethernet1/2
Switch(config-if)# authentication port-control auto
Switch(config-if)# authentication host-mode multi-domain
Switch(config-if)# switchport voice vlan 101
Switch(config-if)# end
```

Disabling the VLAN Upon Switch Violoation: Example

This example shows how to configure the switch to shut down any VLAN on which a security violation error occurs:

```
Switch(config)# errdisable detect cause security-violation shutdown vlan
```

This example shows how to reenable all VLANs that were error-disabled:

```
Switch# clear errdisable interface gigabitethernet1/2 vlan
```

You can verify your settings by entering the **show errdisable detect** privileged EXEC command.

Configuring the Radius Server Parameters: Example

This example shows how to specify the server with IP address 172.20.39.46 as the RADIUS server, to use port 1612 as the authorization port, and to set the encryption key to *rad123*, matching the key on the RADIUS server:

```
Switch(config)# radius-server host 172.120.39.46 auth-port 1612 key rad123
```

Configuring 802.1x Accounting: Example

This example shows how to configure 802.1x accounting. The first command configures the RADIUS server, specifying 1813 as the UDP port for accounting:

```
Switch(config)# radius-server host 172.120.39.46 auth-port 1812 acct-port 1813 key rad123
Switch(config)# aaa accounting dot1x default start-stop group radius
Switch(config)# aaa accounting system default start-stop group radius
```

Enabling an 802.1x Guest VLAN: Example

This example shows how to enable VLAN 2 as an 802.1x guest VLAN:

```
Switch(config)# interface gigabitethernet1/2
Switch(config-if)# authentication event no-response action authorize vlan 2
```

This example shows how to set 3 as the quiet time on the switch, to set 15 as the number of seconds that the switch waits for a response to an EAP-request/identity frame from the client before resending the request, and to enable VLAN 2 as an 802.1x guest VLAN when an 802.1x port is connected to a DHCP client:

```
Switch(config-if)# authentication timer inactivity 3
Switch(config-if)# authentication timer reauthenticate 15
Switch(config-if)# authentication event no-response action authorize vlan 2
```

Displaying Authentication Manager Common Session ID: Examples

This example shows how the session ID appears in the output of the **show authentication** command. The session ID in this example is 160000050000000B288508E5:

```
Switch# show authentication sessions

Interface  MAC Address     Method   Domain   Status         Session ID
Fa4/0/4    0000.0000.0203  mab      DATA     Authz Success  160000050000000B288508E5
```

This is an example of how the session ID appears in the syslog output. The session ID in this example is also 160000050000000B288508E5:

```
1w0d: %AUTHMGR-5-START: Starting 'mab' for client (0000.0000.0203) on Interface Fa4/0/4
AuditSessionID 160000050000000B288508E5
1w0d: %MAB-5-SUCCESS: Authentication successful for client (0000.0000.0203) on Interface
Fa4/0/4 AuditSessionID 160000050000000B288508E5
1w0d: %AUTHMGR-7-RESULT: Authentication result 'success' from 'mab' for client
(0000.0000.0203) on Interface Fa4/0/4 AuditSessionID 160000050000000B288508E5
```

The session ID is used by the NAD, the AAA server, and other report-analyzing applications to identify the client. The ID appears automatically. No configuration is required.

Configuring Inaccessible Authentication Bypass: Example

This example shows how to configure the inaccessible authentication bypass feature:

```
Switch(config)# radius-server dead-criteria time 30 tries 20
Switch(config)# radius-server deadtime 60
Switch(config)# radius-server host 1.1.1.2 acct-port 1550 auth-port 1560 test username
user1 idle-time 30 key abc1234
Switch(config)# dot1x critical eapol
Switch(config)# dot1x critical recovery delay 2000
Switch(config)# interface gigabitethernet 1/1
Switch(config)# radius-server deadtime 60
Switch(config-if)# dot1x critical
Switch(config-if)# dot1x critical recovery action reinitialize
Switch(config-if)# dot1x critical vlan 20
Switch(config-if)# end
```

Configuring VLAN Groups: Examples

This example shows how to configure the VLAN groups, to map the VLANs to the groups, and to verify the VLAN group configurations and mapping to the specified VLANs:

```
switch(config)# vlan group eng-dept vlan-list 10

switch(config)# show vlan group group-name eng-dept
Group Name                      Vlans Mapped
-------------                   --------------
eng-dept                        10
switch# show dot1x vlan-group all
Group Name                      Vlans Mapped
-------------                   --------------
eng-dept                        10
hr-dept                         20
```

This example shows how to add a VLAN to an existing VLAN group and to verify that the VLAN was added:

```
switch(config)# vlan group eng-dept vlan-list 30
switch(config)# show vlan group eng-dept
Group Name                      Vlans Mapped
-------------                   --------------
eng-dept                        10,30
```

This example shows how to remove a VLAN from a VLAN group:

```
switch# no vlan group eng-dept vlan-list 10
```

This example shows that when all the VLANs are cleared from a VLAN group, the VLAN group is cleared:

```
switch(config)# no vlan group eng-dept vlan-list 30
Vlan 30 is successfully cleared from vlan group eng-dept.

switch(config)# show vlan group group-name eng-dept
```

This example shows how to clear all the VLAN groups:

```
switch(config)# no vlan group end-dept vlan-list all
switch(config)# show vlan-group all
```

For more information about these commands, see the *Cisco IOS Security Command Reference.*

Configuring NAC Layer 2 802.1x Validation: Example

This example shows how to configure NAC Layer 2 802.1x validation:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# authentication periodic
Switch(config-if)# authentication timer reauthenticate
```

Configuring an 802.1x Authenticator Switch: Example

This example shows how to configure a switch as an 802.1x authenticator:

```
Switch# configure terminal
Switch(config)# cisp enable
```

```
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# switchport mode access
Switch(config-if)# authentication port-control auto
Switch(config-if)# dot1x pae authenticator
Switch(config-if)# spanning-tree portfast trunk
```

Configuring an 802.1x Supplicant Switch: Example

This example shows how to configure a switch as a supplicant:

```
Switch# configure terminal
Switch(config)# cisp enable
Switch(config)# dot1x credentials test
Switch(config)# username suppswitch
Switch(config)# password myswitch
Switch(config)# dot1x supplicant force-multicast
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport mode trunk
Switch(config-if)# dot1x pae supplicant
Switch(config-if)# dot1x credentials test
Switch(config-if)# end
```

Configuring a Downloadable Policy: Example

This example shows how to configure a switch for a downloadable policy:

```
Switch# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# aaa new-model
Switch(config)# aaa authorization network default group radius
Switch(config)# ip device tracking
Switch(config)# ip access-list extended default_acl
Switch(config-ext-nacl)# permit ip any any
Switch(config-ext-nacl)# exit
Switch(config)# radius-server vsa send authentication
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# ip access-group default_acl in
Switch(config-if)# exit
```

Configuring Open 1x on a Port: Example

This example shows how to configure open 1x on a port:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/1
Switch(config)# authentication control-direction both
Switch(config)# au ten tic at ion fallback profile1
Switch(config)# authentication host-mode multi-auth
Switch(config)# authentication open
Switch(config)# authentication order dot1x webauth
Switch(config)# authentication periodic
Switch(config)# authentication port-control auto
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
Radius commands	*Cisco IOS Security Command Reference*
Switch authentication configuration	Chapter 12, “Configuring Switch-Based Authentication”
Authenticator switch information	Chapter 16, “Configuring Smartports Macros”

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 14

Configuring Web-Based Authentication

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for Configuring Web-Based Authentication

- By default, the IP device tracking feature is disabled on a switch. You must enable the IP device tracking feature to use web-based authentication.
- You must configure at least one IP address to run the switch HTTP server. You must also configure routes to reach each host IP address. The HTTP server sends the HTTP login page to the host.
- You must configure the default ACL on the interface before configuring web-based authentication. Configure a port ACL for a Layer 2 interface.

Restrictions for Configuring Web-Based Authentication on the IE 2000 Switch

- Web-based authentication is an ingress-only feature.
- You can configure web-based authentication only on access ports. Web-based authentication is not supported on trunk ports, EtherChannel member ports, or dynamic trunk ports.
- You cannot authenticate hosts on Layer 2 interfaces with static ARP cache assignment. These hosts are not detected by the web-based authentication feature because they do not send ARP messages.
- Hosts that are more than one hop away might experience traffic disruption if an STP topology change results in the host traffic arriving on a different port. This occurs because the ARP and DHCP updates might not be sent after a Layer 2 (STP) topology change.
- Web-based authentication does not support VLAN assignment as a downloadable-host policy.
- Web-based authentication is not supported for IPv6 traffic.

- Web-based authentication and Network Edge Access Topology (NEAT) are mutually exclusive. You cannot use web-based authentication when NEAT is enabled on an interface, and you cannot use NEAT when web-based authentication is running on an interface.
- Web-based authentication supports only RADIUS authorization servers. You cannot use TACACS+ servers or local authorization.

Information About Configuring Web-Based Authentication

Web-Based Authentication

Use the web-based authentication feature, known as *web authentication proxy*, to authenticate end users on host systems that do not run the IEEE 802.1x supplicant.

Note You can configure web-based authentication on Layer 2 interfaces.

When you initiate an HTTP session, web-based authentication intercepts ingress HTTP packets from the host and sends an HTML login page to the users. The users enter their credentials, which the web-based authentication feature sends to the authentication, authorization, and accounting (AAA) server for authentication.

If authentication succeeds, web-based authentication sends a Login-Successful HTML page to the host and applies the access policies returned by the AAA server.

If authentication fails, web-based authentication forwards a Login-Fail HTML page to the user, prompting the user to retry the login. If the user exceeds the maximum number of attempts, web-based authentication forwards a Login-Expired HTML page to the host, and the user is placed on a watch list for a waiting period.

These sections describe the role of web-based authentication as part of AAA:

- Device Roles, page 14-2
- Host Detection, page 14-3
- Session Creation, page 14-3
- Authentication Process, page 14-4
- Web Authentication Customizable Web Pages, page 14-6
- Web-Based Authentication Interactions with Other Features, page 14-8

Device Roles

With web-based authentication, the devices in the network have these specific roles:

- Client—The device (workstation) that requests access to the LAN and the services and responds to requests from the switch. The workstation must be running an HTML browser with Java Script enabled.
- Authentication server—Authenticates the client. The authentication server validates the identity of the client and notifies the switch that the client is authorized to access the LAN and the switch services or that the client is denied.

- Switch—Controls the physical access to the network based on the authentication status of the client. The switch acts as an intermediary (proxy) between the client and the authentication server, requesting identity information from the client, verifying that information with the authentication server, and relaying a response to the client.

Figure 14-1 Web-Based Authentication Device Roles

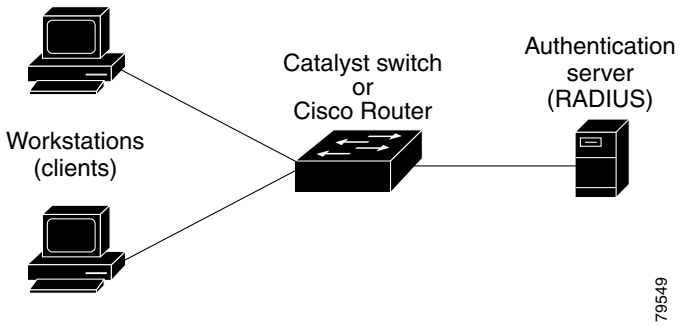

Host Detection

The switch maintains an IP device tracking table to store information about detected hosts.

Note By default, the IP device tracking feature is disabled on a switch. You must enable the IP device tracking feature to use web-based authentication.

For Layer 2 interfaces, web-based authentication detects IP hosts by using these mechanisms:

- ARP-based trigger—ARP redirect ACL allows web-based authentication to detect hosts with a static IP address or a dynamic IP address.
- Dynamic ARP inspection
- DHCP snooping—Web-based authentication is notified when the switch creates a DHCP-binding entry for the host.

Session Creation

When web-based authentication detects a new host, it creates a session as follows:

- Reviews the exception list.

  If the host IP is included in the exception list, the policy from the exception list entry is applied, and the session is established.

- Reviews for authorization bypass.

  If the host IP is not on the exception list, web-based authentication sends a nonresponsive-host (NRH) request to the server.

  If the server response is *access accepted*, authorization is bypassed for this host. The session is established.

- Sets up the HTTP intercept ACL.

If the server response to the NRH request is *access rejected*, the HTTP intercept ACL is activated, and the session waits for HTTP traffic from the host.

Authentication Process

When you enable web-based authentication, these events occur:

- The user initiates an HTTP session.
- The HTTP traffic is intercepted, and authorization is initiated. The switch sends the login page to the user. The user enters a username and password, and the switch sends the entries to the authentication server.
- If the authentication succeeds, the switch downloads and activates the user's access policy from the authentication server. The login success page is sent to the user.
- If the authentication fails, the switch sends the login fail page. The user retries the login. If the maximum number of attempts fails, the switch sends the login expired page, and the host is placed in a watch list. After the watch list times out, the user can retry the authentication process.
- If the authentication server does not respond to the switch, and if an AAA fail policy is configured, the switch applies the failure access policy to the host. The login success page is sent to the user. (See the "Local Web Authentication Banner" section on page 14-4.)
- The switch reauthenticates a client when the host does not respond to an ARP probe on a Layer 2 interface, or when the host does not send any traffic within the idle timeout on a Layer 3 interface.
- The feature applies the downloaded timeout or the locally configured session timeout.
- If the terminate action is RADIUS, the feature sends a nonresponsive host (NRH) request to the server. The terminate action is included in the response from the server.
- If the terminate action is default, the session is dismantled, and the applied policy is removed.

Local Web Authentication Banner

You can create a banner that will appear when you log in to a switch by using web authentication.

The banner appears on both the login page and the authentication-result pop-up pages:

- Authentication Successful
- Authentication Failed
- Authentication Expired

You create a banner by using the **ip admission auth-proxy-banner http** global configuration command. The default banner Cisco Systems and Switch host-name Authentication appear on the Login Page. Cisco Systems appears on the authentication result pop-up page, as shown in Figure 14-2.

Figure 14-2 Authentication Successful Banner

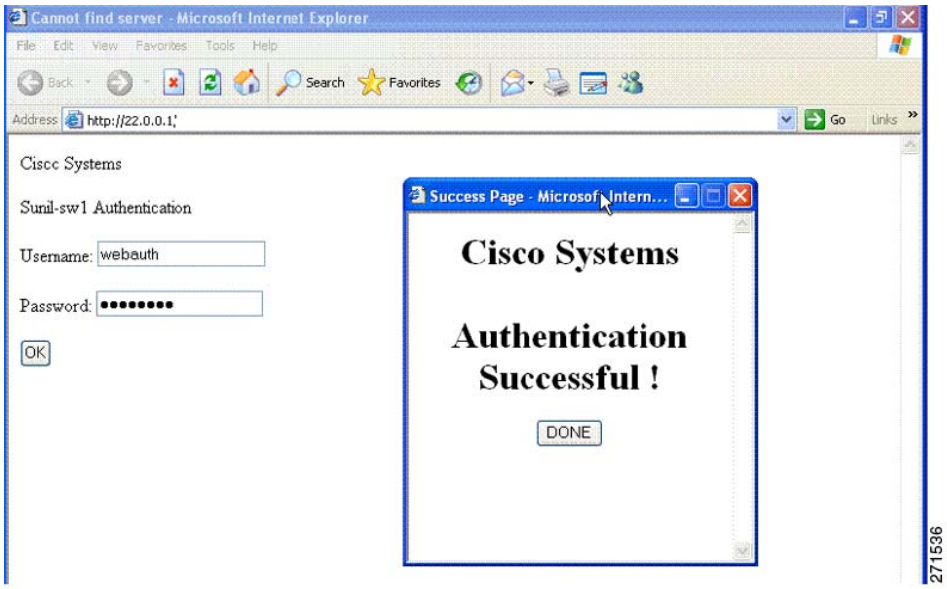

You can also customize the banner, as shown in Figure 14-3.

- Add a switch, router, or company name to the banner by using the **ip admission auth-proxy-banner http** *banner-text* global configuration command.
- Add a logo or text file to the banner by using the **ip admission auth-proxy-banner http** *file-path* global configuration command.

Figure 14-3 Customized Web Banner

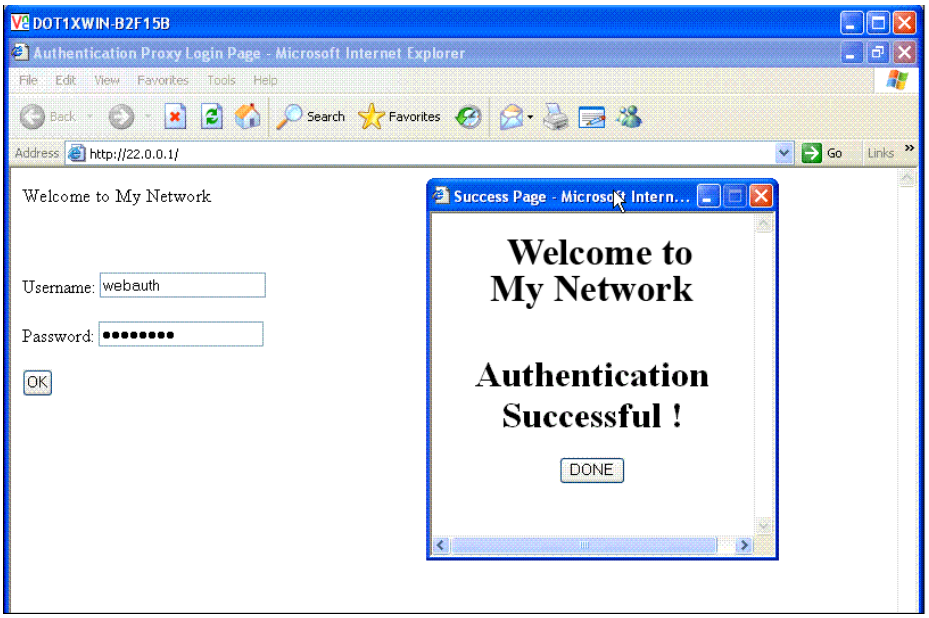

If you do not enable a banner, only the username and password dialog boxes appear in the web authentication login screen, and no banner appears when you log into the switch, as shown in Figure 14-4.

Figure 14-4 Login Screen with No Banner

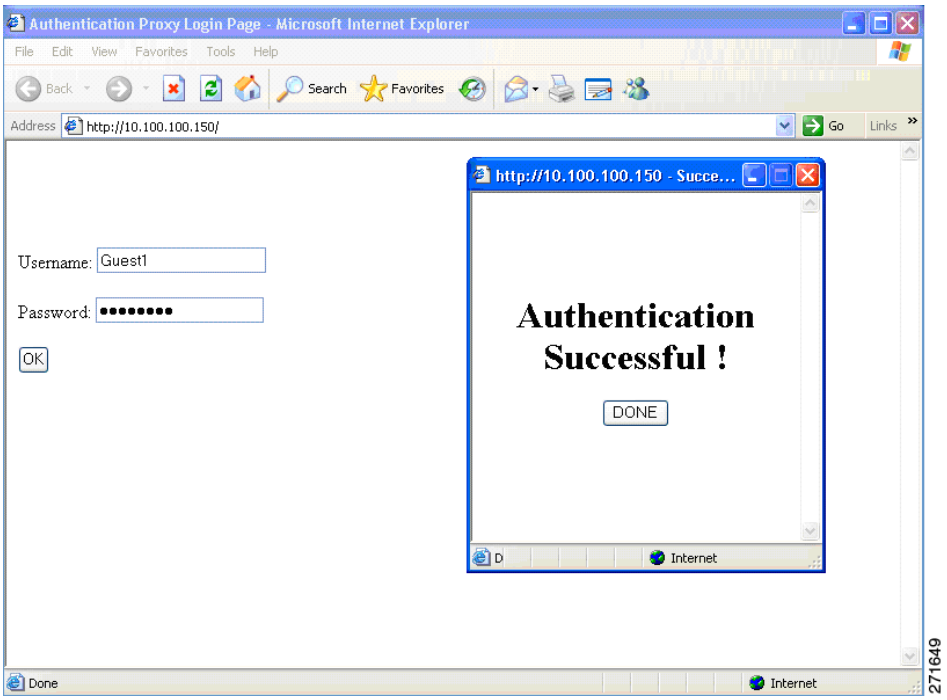

For more information, see the *Cisco IOS Security Command Reference* and the "Configuring a Web Authentication Local Banner" section on page 14-14.

Web Authentication Customizable Web Pages

During the web-based authentication process, the switch internal HTTP server hosts four HTML pages to deliver to an authenticating client. The server uses these pages to notify you of these four-authentication process states:

- Login—Your credentials are requested.
- Success—The login was successful.
- Fail—The login failed.
- Expire—The login session has expired because of excessive login failures.

Web Authentication Guidelines

- You can substitute your own HTML pages for the default internal HTML pages.
- You can use a logo or specify text in the login, success, failure, and expire web pages.
- On the banner page, you can specify text in the login page.
- The pages are in HTML.
- You must include an HTML redirect command in the success page to access a specific URL.
- The URL string must be a valid URL (for example, http://www.cisco.com). An incomplete URL might cause page not found error or similar errors on a web browser.

- If you configure web pages for HTTP authentication, they must include the appropriate HTML commands (for example, to set the page time out, to set a hidden password, or to confirm that the same page is not submitted twice).
- The CLI command to redirect users to a specific URL is not available when the configured login form is enabled. The administrator should ensure that the redirection is configured in the web page.
- If the CLI command redirecting users to a specific URL after authentication occurs is entered and then the command configuring web pages is entered, the CLI command redirecting users to a specific URL does not take effect.
- Configured web pages can be copied to the switch boot flash or flash.
- Configured pages can be accessed from the flash on the stack master or members.
- The login page can be on one flash, and the success and failure pages can be another flash (for example, the flash on the stack master or a member).
- You must configure all four pages.
- The banner page has no effect if it is configured with the web page.
- All of the logo files (image, flash, audio, video, and so on) that are stored in the system directory (for example, flash, disk0, or disk) and that must be displayed on the login page must use web_auth_*filename* as the filename.
- The configured authentication proxy feature supports both HTTP and SSL.

When configuring customized authentication proxy web pages, follow these guidelines:

- To enable the custom web pages feature, specify all four custom HTML files. If you specify fewer than four files, the internal default HTML pages are used.
- The four custom HTML files must be present on the flash memory of the switch. The maximum size of each HTML file is 8 KB.
- Any images on the custom pages must be on an accessible HTTP server. Configure an intercept ACL within the admission rule.
- Any external link from a custom page requires configuration of an intercept ACL within the admission rule.
- To access a valid DNS server, any name resolution required for external links or images requires configuration of an intercept ACL within the admission rule.
- If the custom web pages feature is enabled, a configured auth-proxy-banner is not used.
- If the custom web pages feature is enabled, the redirection URL for successful login feature is not available.
- To remove the specification of a custom file, use the **no** form of the command.

Because the custom login page is a public web form, consider these guidelines for the page:

- The login form must accept user entries for the username and password and must show them as **uname** and **pwd**.
- The custom login page should follow best practices for a web form, such as page timeout, hidden password, and prevention of redundant submissions.

You can substitute your HTML pages, as shown in Figure 14-5, for the default internal HTML pages. You can also specify a URL to which users are redirected after authentication occurs, which replaces the internal Success page.

Figure 14-5 Customizeable Authentication Page

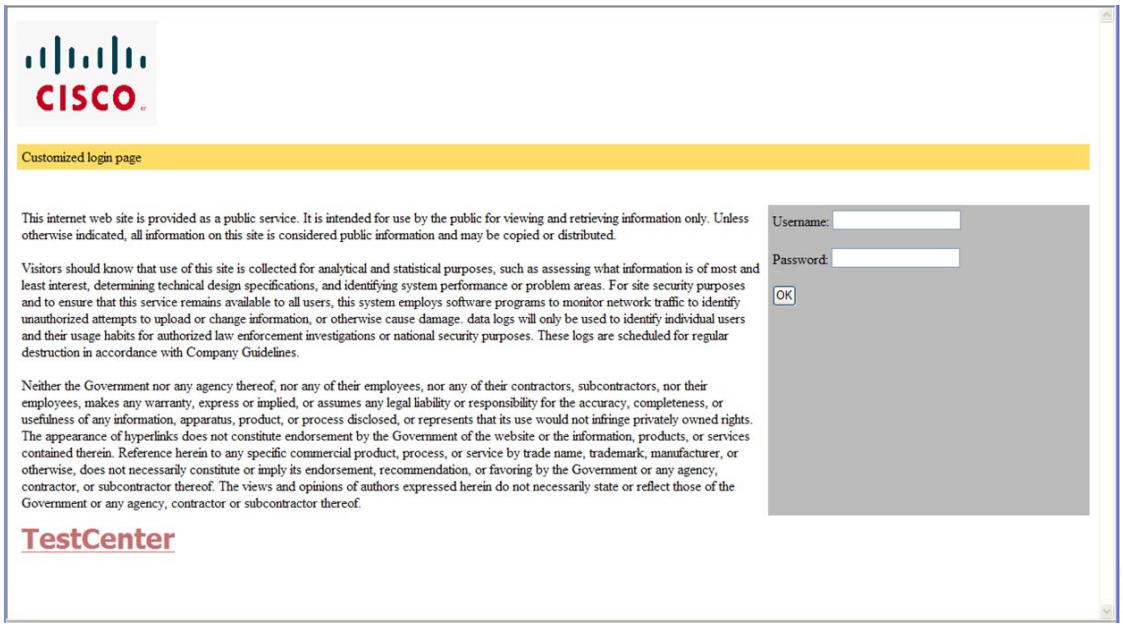

Web-Based Authentication Interactions with Other Features

- Port Security, page 14-8
- LAN Port IP, page 14-8
- Gateway IP, page 14-9
- ACLs, page 14-9
- Context-Based Access Control, page 14-9
- 802.1x Authentication, page 14-9
- EtherChannel, page 14-9

Port Security

You can configure web-based authentication and port security on the same port. Web-based authentication authenticates the port, and port security manages network access for all MAC addresses, including that of the client. You can then limit the number or group of clients that can access the network through the port.

For more information about enabling port security, see the "Configuring Port Security" section on page 29-11.

LAN Port IP

You can configure LAN port IP (LPIP) and Layer 2 web-based authentication on the same port. The host is authenticated by using web-based authentication first, followed by LPIP posture validation. The LPIP host policy overrides the web-based authentication host policy.

If the web-based authentication idle timer expires, the NAC policy is removed. The host is authenticated, and posture is validated again.

Gateway IP

You cannot configure Gateway IP (GWIP) on a Layer 3 VLAN interface if web-based authentication is configured on any of the switch ports in the VLAN.

You can configure web-based authentication on the same Layer 3 interface as Gateway IP. The host policies for both features are applied in software. The GWIP policy overrides the web-based authentication host policy.

ACLs

If you configure a VLAN ACL or a Cisco IOS ACL on an interface, the ACL is applied to the host traffic only after the web-based authentication host policy is applied.

For Layer 2 web-based authentication, you must configure a port ACL (PACL) as the default access policy for ingress traffic from hosts connected to the port. After authentication, the web-based authentication host policy overrides the PACL.

Note When a proxy ACL is configured for a web-based authentication client, the proxy ACL is downloaded and applied as part of the authorization process. Hence, the PACL displays the proxy ACL access control entry (ACE).

You cannot configure a MAC ACL and web-based authentication on the same interface.

You cannot configure web-based authentication on a port whose access VLAN is configured for VACL capture.

Context-Based Access Control

Web-based authentication cannot be configured on a Layer 2 port if context-based access control (CBAC) is configured on the Layer 3 VLAN interface of the port VLAN.

802.1x Authentication

You cannot configure web-based authentication on the same port as 802.1x authentication except as a fallback authentication method.

EtherChannel

You can configure web-based authentication on a Layer 2 EtherChannel interface. The web-based authentication configuration applies to all member channels.

Default Web-Based Authentication Settings

Table 14-1 Default Web-Based Authentication Settings

Feature	Default Settings
AAA	Disabled
RADIUS server • IP address • UDP authentication port • Key	 • None specified • 1812 • None specified
Default value of inactivity timeout	3600 seconds
Inactivity timeout	Enabled

Configuring Switch-to-RADIUS-Server Communication

RADIUS security servers identification:

- Host name
- Host IP address
- Host name and specific UDP port numbers
- IP address and specific UDP port numbers

The combination of the IP address and UDP port number creates a unique identifier, that enables RADIUS requests to be sent to multiple UDP ports on a server at the same IP address. If two different host entries on the same RADIUS server are configured for the same service (for example, authentication) the second host entry that is configured functions as the failover backup to the first one. The RADIUS host entries are chosen in the order that they were configured.

How to Configure Web-Based Authentication

Configuring the Authentication Rule and Interfaces

	Command	Purpose
Step 1	**ip admission name** *name* **proxy http**	Configures an authentication rule for web-based authorization.
Step 2	**interface** *type slot/port*	Enters interface configuration mode and specifies the ingress Layer 2 interface to be enabled for web-based authentication. *type* can be Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet.
Step 3	**ip access-group** *name*	Applies the default ACL.
Step 4	**ip admission** *name*	Configures web-based authentication on the specified interface.
Step 5	**exit**	Returns to configuration mode.
Step 6	**ip device tracking**	Enables the IP device tracking table.
Step 7	**end**	Returns to privileged EXEC mode.
Step 8	**show ip admission configuration**	Displays the configuration.

Configuring AAA Authentication

	Command	Purpose
Step 1	**aaa new-model**	Enables AAA functionality.
Step 2	**aaa authentication login default group** {*tacacs+* | *radius*}	Defines the list of authentication methods at login.
Step 3	**aaa authorization auth-proxy default group** {*tacacs+* | *radius*}	Creates an authorization method list for web-based authorization.
Step 4	**radius-server host** {*hostname* | *ip-address*} **test username** *username*	Specifies an AAA server. Specifies the host name or IP address of the remote RADIUS server. The **test username** *username* option enables automated testing of the RADIUS server connection. The specified username does not need to be a valid user name.
Step 5	**radius-server key** *string*	Configures the authorization and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. To use multiple RADIUS servers, reenter this command for each server.

Configuring Switch-to-RADIUS-Server Communication

	Command	Purpose
Step 1	**ip radius source-interface** *interface_name*	Specifies that the RADIUS packets have the IP address of the indicated interface.
Step 2	**radius-server host** {*hostname* | *ip-address*} **test username** *username*	Specifies the host name or IP address of the remote RADIUS server. The **test username** *username* option enables automated testing of the RADIUS server connection. The specified *username* does not need to be a valid user name. The **key** option specifies an authentication and encryption key to use between the switch and the RADIUS server. To use multiple RADIUS servers, reenter this command for each server.
Step 3	**radius-server key** *string*	Configures the authorization and encryption key used between the switch and the RADIUS daemon running on the RADIUS server.
Step 4	**radius-server vsa send authentication**	Enables downloading of an ACL from the RADIUS server. This feature is supported in Cisco IOS Release 12.2(50)SG.
Step 5	**radius-server dead-criteria tries** *num-tries*	Specifies the number of unanswered sent messages to a RADIUS server before considering the server to be inactive. The range of *num-tries* is 1 to 100.

Configuring the HTTP Server

	Command	Purpose
Step 1	**ip http server**	Enables the HTTP server. The web-based authentication feature uses the HTTP server to communicate with the hosts for user authentication.
Step 2	**ip http secure-server**	Enables HTTPS.

Customizing the Authentication Proxy Web Pages

Before You Begin

You can configure web authentication to display four substitute HTML pages to the user in place of the switch default HTML pages during web-based authentication.

To specify the use of your custom authentication proxy web pages, first store your custom HTML files on the switch flash memory, then perform this task in global configuration mode:

	Command	Purpose
Step 1	**ip admission proxy http login page file** *device:login-filename*	Specifies the location in the switch memory file system of the custom HTML file to use in place of the default login page. The *device:* is flash memory.
Step 2	**ip admission proxy http success page file** *device:success-filename*	Specifies the location of the custom HTML file to use in place of the default login success page.
Step 3	**ip admission proxy http failure page file** *device:fail-filename*	Specifies the location of the custom HTML file to use in place of the default login failure page.
Step 4	**ip admission proxy http login expired page file** *device:expired-filename*	Specifies the location of the custom HTML file to use in place of the default login expired page.

Specifying a Redirection URL for Successful Login

You can specify a URL to which the user is redirected after authentication, effectively replacing the internal S*uccess* HTML page.

Command	Purpose
ip admission proxy http success redirect *url-string*	Specifies a URL for redirection of the user in place of the default login success page.

Configuring the Web-Based Authentication Parameters

You can configure the maximum number of failed login attempts before the client is placed in a watch list for a waiting period.

	Command	Purpose
Step 1	**ip admission max-login-attempts** *number*	Sets the maximum number of failed login attempts. The range is 1 to 2147483647 attempts. The default is 5.
Step 2	**end**	Returns to privileged EXEC mode.
Step 3	**show ip admission configuration**	Displays the authentication proxy configuration.
Step 4	**show ip admission cache**	Displays the list of authentication entries.
Step 5	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring a Web Authentication Local Banner

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip admission auth-proxy-banner http** [*banner-text* | *file-path*]	Enables the local banner. (Optional) Creates a custom banner by entering *C banner-text C,* where *C* is a delimiting character or a file-path indicates a file (for example, a logo or text file) that appears in the banner.
Step 3	**end**	Returns to privileged EXEC mode.
Step 4	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Removing Web-Based Authentication Cache Entries

Enter a specific IP address to delete the entry for a single host. Use an asterisk to delete all cache entries.

Command	Purpose
clear ip auth-proxy cache {* | *host ip address*}	Clears authentication proxy entries from the switch.
clear ip admission cache {* | *host ip address*}	Clears IP admission cache entries from the switch.

Monitoring and Maintaining Web-Based Authentication

Command	Purpose
show authentication sessions	Displays the web-based authentication settings.
show ip admission configuration	Displays the authentication proxy configuration.
show ip admission cache	Displays the list of authentication entries.

Configuration Examples for Configuring Web-Based Authentication

Enabling and Displaying Web-Based Authentication: Examples

This example shows how to enable web-based authentication on Fast Ethernet port 5/1:

```
Switch(config)# ip admission name webauth1 proxy http
Switch(config)# interface fastethernet 5/1
Switch(config-if)# ip admission webauth1
Switch(config-if)# exit
Switch(config)# ip device tracking
```

This example shows how to verify the configuration:

```
Switch# show ip admission configuration
Authentication Proxy Banner not configured
Authentication global cache time is 60 minutes
Authentication global absolute time is 0 minutes
Authentication global init state time is 2 minutes
Authentication Proxy Watch-list is disabled

Authentication Proxy Rule Configuration
 Auth-proxy name webauth1
    http list not specified inactivity-time 60 minutes

Authentication Proxy Auditing is disabled
Max Login attempts per user is 5
```

Enabling AAA: Example

This example shows how to enable AAA:

```
Switch(config)# aaa new-model
Switch(config)# aaa authentication login default group radius
Switch(config)# aaa authorization auth-proxy default group radius
```

Configuring the RADIUS Server Parameters: Example

This example shows how to configure the RADIUS server parameters on a switch:

```
Switch(config)# ip radius source-interface Vlan80
Switch(config)# radius-server host 172.120.39.46 test username user1
Switch(config)# radius-server key rad123
Switch(config)# radius-server dead-criteria tries 2
```

Configuring a Custom Authentication Proxy Web Page: Example

This example shows how to configure custom authentication proxy web pages:

```
Switch(config)# ip admission proxy http login page file flash:login.htm
Switch(config)# ip admission proxy http success page file flash:success.htm
Switch(config)# ip admission proxy http fail page file flash:fail.htm
Switch(config)# ip admission proxy http login expired page flash flash:expired.htm
```

Verifying a Custom Authentication Proxy Web Page: Example

This example shows how to verify the configuration of a custom authentication proxy web pages:

```
Switch# show ip admission configuration
Authentication proxy webpage
 Login page           : flash:login.htm
 Success page         : flash:success.htm
 Fail Page            : flash:fail.htm
 Login expired Page   : flash:expired.htm

Authentication global cache time is 60 minutes
Authentication global absolute time is 0 minutes
Authentication global init state time is 2 minutes
Authentication Proxy Session ratelimit is 100
```

```
Authentication Proxy Watch-list is disabled
Authentication Proxy Auditing is disabled
Max Login attempts per user is 5
```

Configuring a Redirection URL: Example

This example shows how to configure a redirection URL for successful login:

```
Switch(config)# ip admission proxy http success redirect www.cisco.com
```

Verifying a Redirection URL: Example

This example shows how to verify the redirection URL for successful login:

```
Switch# show ip admission configuration
Authentication Proxy Banner not configured
Customizable Authentication Proxy webpage not configured
HTTP Authentication success redirect to URL: http://www.cisco.com
Authentication global cache time is 60 minutes
Authentication global absolute time is 0 minutes
Authentication global init state time is 2 minutes
Authentication Proxy Watch-list is disabled
Authentication Proxy Max HTTP process is 7
Authentication Proxy Auditing is disabled
Max Login attempts per user is 5
```

Configuring a Local Banner: Example

This example shows how to configure a local banner with the custom message *My Switch*:

```
Switch(config) configure terminal
Switch(config)# aaa new-model
Switch(config)# aaa ip auth-proxy auth-proxy-banner C My Switch C
Switch(config) end
```

Clearing the Web-Based Authentication Session: Example

This example shows how to remove the web-based authentication session for the client at the IP address 209.165.201.1:

```
Switch# clear ip auth-proxy cache 209.165.201.1
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
Authentication proxy commands Radius server commands	*Cisco IOS Security Command Reference*
Authentication proxy configuration Radius server configuration	*Cisco IOS Security Configuration Guide*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 15

Configuring Interface Characteristics

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Restrictions for Configuring Interface Characteristics

- The EtherChannel port group interface is supported on a switch running the LAN Base image.

Information About Configuring Interface Characteristics

Interface Types

This section describes the different types of interfaces supported by the switch with references to chapters that contain more detailed information about configuring these interface types.

- Port-Based VLANs, page 15-2
- Switch Ports, page 15-2
- Routed Ports, page 15-3
- EtherChannel Port Groups, page 15-4
- Dual-Purpose Uplink Ports, page 15-4
- Connecting Interfaces, page 15-5

Port-Based VLANs

A VLAN is a switched network that is logically segmented by function, team, or application, without regard to the physical location of the users. For more information about VLANs, see the Chapter 17, "Configuring VLANs." Packets received on a port are forwarded only to ports that belong to the same VLAN as the receiving port. Network devices in different VLANs cannot communicate with one another without a Layer 3 device to route traffic between the VLANs.

VLAN partitions provide hard firewalls for traffic in the VLAN, and each VLAN has its own MAC address table. A VLAN comes into existence when a local port is configured to be associated with the VLAN, when the VLAN Trunking Protocol (VTP) learns of its existence from a neighbor on a trunk, or when a user creates a VLAN.

To configure VLANs, use the **vlan** *vlan-id* global configuration command to enter VLAN configuration mode. The VLAN configurations for normal-range VLANs (VLAN IDs 1 to 1005) are saved in the VLAN database. If VTP is version 1 or 2, to configure extended-range VLANs (VLAN IDs 1006 to 4096), you must first set VTP mode to transparent. Extended-range VLANs created in transparent mode are not added to the VLAN database but are saved in the switch running configuration. With VTP version 3, you can create extended-range VLANs in client or server mode. These VLANs are saved in the VLAN database.

Add ports to a VLAN by using the **switchport** interface configuration commands:

- Identify the interface.
- For a trunk port, set trunk characteristics, and if desired, define the VLANs to which it can belong.
- For an access port, set and define the VLAN to which it belongs.

Switch Ports

Switch ports are Layer 2-only interfaces associated with a physical port. A switch port can be an access port, a trunk port, or a tunnel port. You can configure a port as an access port or trunk port or let the Dynamic Trunking Protocol (DTP) operate on a per-port basis to set the switchport mode by negotiating with the port on the other end of the link. Switch ports are used for managing the physical interface and associated Layer 2 protocols.

Configure switch ports by using the **switchport** interface configuration commands. Use the **switchport** command with no keywords to put an interface that is in Layer 3 mode into Layer 2 mode.

Note When you put an interface that is in Layer 3 mode into Layer 2 mode, the previous configuration information related to the affected interface might be lost, and the interface is returned to its default configuration.

For detailed information about configuring access port and trunk port characteristics, see Chapter 17, "Configuring VLANs."

Routed Ports

Note The LAN base image supports static routing.

A routed port is a physical port that acts like a port on a router; it does not have to be connected to a router. A routed port is not associated with a particular VLAN, as is an access port. A routed port behaves like a regular router interface, except that it does not support VLAN subinterfaces. Routed ports can be configured with a Layer 3 routing protocol. A routed port is a Layer 3 interface only and does not support Layer 2 protocols, such as DTP and STP. Routed ports are supported only on switches running the IP base or IP services image.

Configure routed ports by putting the interface into Layer 3 mode with the **no switchport** interface configuration command. Then assign an IP address to the port, enable routing, and assign routing protocol characteristics by using the **ip routing** and **router** *protocol* global configuration commands.

Note Entering a **no switchport** interface configuration command shuts down the interface and then reenables it, which might generate messages on the device to which the interface is connected. When you put an interface that is in Layer 2 mode into Layer 3 mode, the previous configuration information related to the affected interface might be lost.

The number of routed ports that you can configure is not limited by software. However, the interrelationship between this number and the number of other features being configured might impact CPU performance because of hardware limitations.

For more information about IP unicast routing and routing protocols, see Chapter 41, “Configuring IP Unicast Routing”

Access Ports

An access port belongs to and carries the traffic of only one VLAN (unless it is configured as a voice VLAN port). Traffic is received and sent in native formats with no VLAN tagging. Traffic arriving on an access port is assumed to belong to the VLAN assigned to the port.

If an access port receives an 802.1Q tagged packet, the packet is dropped, and the source address is not learned.

Two types of access ports are supported:

- Static access ports are manually assigned to a VLAN (or through a RADIUS server for use with IEEE 802.1x. For more information, see the “802.1x Authentication with VLAN Assignment” section on page 13-15.
- VLAN membership of dynamic access ports is learned through incoming packets. By default, a dynamic access port is not a member of any VLAN, and forwarding to and from the port is enabled only when the VLAN membership of the port is discovered. Dynamic access ports on the switch are assigned to a VLAN by a VLAN Membership Policy Server (VMPS). The VMPS can be a Catalyst 6500 series switch; the switch cannot be a VMPS server.

You can also configure an access port with an attached Cisco IP Phone to use one VLAN for voice traffic and another VLAN for data traffic from a device attached to the phone. For more information about voice VLAN ports, see Chapter 19, “Configuring Voice VLAN.”

Trunk Ports

A trunk port carries the traffic of multiple VLANs and by default is a member of all VLANs in the VLAN database.

The switch supports only IEEE 802.1Q trunk ports. An IEEE 802.1Q trunk port supports simultaneous tagged and untagged traffic. An IEEE 802.1Q trunk port is assigned a default port VLAN ID (PVID), and all untagged traffic travels on the port default PVID. All untagged traffic and tagged traffic with a NULL VLAN ID are assumed to belong to the port default PVID. A packet with a VLAN ID equal to the outgoing port default PVID is sent untagged. All other traffic is sent with a VLAN tag.

Although by default, a trunk port is a member of every VLAN known to the VTP, you can limit VLAN membership by configuring an allowed list of VLANs for each trunk port. The list of allowed VLANs does not affect any other port but the associated trunk port. By default, all possible VLANs (VLAN ID 1 to 4096) are in the allowed list. A trunk port can become a member of a VLAN only if VTP knows of the VLAN and if the VLAN is in the enabled state. If VTP learns of a new, enabled VLAN and the VLAN is in the allowed list for a trunk port, the trunk port automatically becomes a member of that VLAN and traffic is forwarded to and from the trunk port for that VLAN. If VTP learns of a new, enabled VLAN that is not in the allowed list for a trunk port, the port does not become a member of the VLAN, and no traffic for the VLAN is forwarded to or from the port.

For more information about trunk ports, see Chapter 17, “Configuring VLANs.”

EtherChannel Port Groups

Note The LAN Base image supports EtherChannel port groups.

EtherChannel port groups treat multiple switch ports as one switch port. These port groups act as a single logical port for high-bandwidth connections between switches or between switches and servers. An EtherChannel balances the traffic load across the links in the channel. If a link within the EtherChannel fails, traffic previously carried over the failed link changes to the remaining links. You can group multiple trunk ports into one logical trunk port, group multiple access ports into one logical access port, group multiple tunnel ports into one logical tunnel port, or group multiple routed ports into one logical routed port.

Most protocols operate over either single ports or aggregated switch ports and do not recognize the physical ports within the port group. Exceptions are the DTP, the Cisco Discovery Protocol (CDP), and the Port Aggregation Protocol (PAgP), which operate only on physical ports.

When you configure an EtherChannel, you create a port-channel logical interface and assign an interface to the EtherChannel. Use the **channel-group** interface configuration command to dynamically create the port-channel logical interface. This command binds the physical and logical ports together.

For Layer 3 interfaces, you manually create the logical interface by using the **interface port-channel** global configuration command. Then you manually assign an interface to the EtherChannel by using the **channel-group** interface configuration command.

For more information, see Chapter 40, “Configuring EtherChannels.”

Dual-Purpose Uplink Ports

Some switches support dual-purpose uplink ports. Each uplink port is considered as a single interface with dual front ends—an RJ-45 connector and a small form-factor pluggable (SFP) module connector. The dual front ends are not redundant interfaces, and the switch activates only one connector of the pair.

By default, the switch dynamically selects the interface type that first links up. However, you can use the **media-type** interface configuration command to manually select the RJ-45 connector or the SFP module connector. To return to the default setting, use the **media-type auto interface** or the **no media-type** interface configuration commands.

Each uplink port has two LEDs: one shows the status of the RJ-45 port, and one shows the status of the SFP module port. The port LED is on for whichever connector is active. For more information about the LEDs, see the hardware installation guide.

The switch configures both types to autonegotiate speed and duplex (the default). If you configure **auto-select**, you cannot configure the **speed** and **duplex** interface configuration commands.

When the switch powers on or when you enable a dual-purpose uplink port through the **shutdown** and the **no shutdown** interface configuration commands, the switch gives preference to the SFP module interface. In all other situations, the switch selects the active link based on which type first links up.

The switch operates with 100BASE-*x* (where -*x* is -BX, -FX-FE, -LX) SFP modules as follows:

- When the 100BASE -*x* SFP module is inserted into the module slot and there is no link on the RJ-45 side, the switch disables the RJ-45 interface and selects the SFP module interface. This is the behavior even if there is no cable connected and if there is no link on the SFP module side.
- When the 100BASE-*x* SFP module is inserted and there is a link on the RJ-45 side, the switch continues with that link. If the link goes down, the switch disables the RJ-45 side and selects the SFP module interface.
- When the 100BASE-*x* SFP module is removed, the switch again dynamically selects the type (**auto-select**) and re-enables the RJ-45 side.

The switch does not have this behavior with 100BASE-FX-GE SFP modules.

Connecting Interfaces

Devices within a single VLAN can communicate directly through any switch. Ports in different VLANs cannot exchange data without going through a routing device.

With a standard Layer 2 switch, ports in different VLANs have to exchange information through a router. By using the switch with routing enabled, when you configure both VLAN 20 and VLAN 30 with an SVI to which an IP address is assigned, packets can be sent from Host A to Host B directly through the switch with no need for an external router (Figure 15-1).

Figure 15-1 Connecting VLANs with a Layer 3 Switch

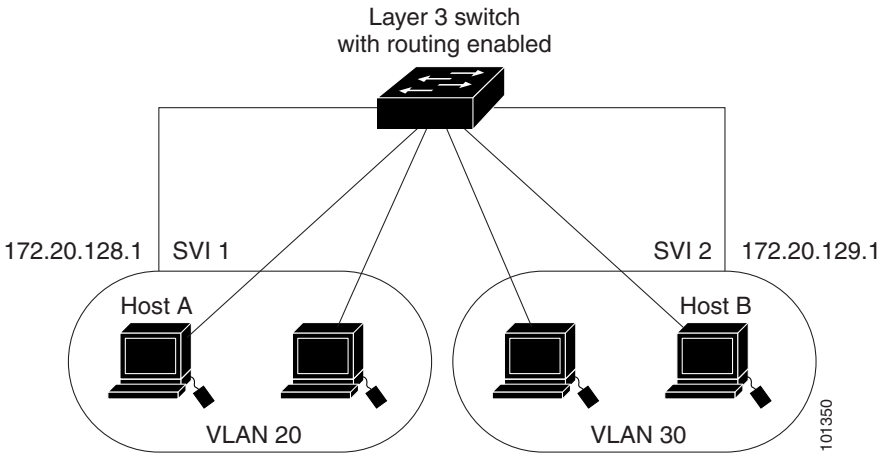

Basic routing (static routing and RIP) is supported on the LAN base image. Whenever possible, to maintain high performance, forwarding is done by the switch hardware. However, only IP Version 4 packets with Ethernet II encapsulation can be routed in hardware. *Non-IP traffic and traffic with other encapsulation methods can be fallback-bridged by hardware.*

The routing function can be enabled on all SVIs and routed ports. The switch routes only IP traffic. When IP routing protocol parameters and address configuration are added to an SVI or routed port, any IP traffic received from these ports is routed.For more information, see Chapter 41, "Configuring IP Unicast Routing."

- *Fallback bridging forwards traffic that the switch does not route or traffic belonging to a nonroutable protocol, such as DECnet. Fallback bridging connects multiple VLANs into one bridge domain by bridging between two or more SVIs or routed ports. When configuring fallback bridging, you assign SVIs or routed ports to bridge groups with each SVI or routed port assigned to only one bridge group. All interfaces in the same group belong to the same bridge domain.*

Using Interface Configuration Mode

The switch supports these interface types:

- Physical ports—switch ports and routed ports
- VLANs—switch virtual interfaces
- Port channels—EtherChannel interfaces

You can also configure a range of interfaces (see the "Configuring a Range of Interfaces" section on page 15-13).

To configure a physical interface (port), specify the interface type, and switch port number, and enter interface configuration mode.

- Type—Port types depend on those supported on the switch. Possible types are Fast Ethernet (fastethernet or fa) for 10/100 Mb/s Ethernet, Gigabit Ethernet (gigabitethernet or gi) for 10/100/1000 Mb/s Ethernet ports, or small form-factor pluggable (SFP) module Gigabit Ethernet interfaces.

- Port number—The physical interface number on the switch. The port numbers for the IE-2000-4TC switch model are 1–4 for the Fast Ethernet ports and 1–2 for the Gigabit Ethernet ports. The port numbers for the IE-2000-8TC switch model are 1–8 for the Fast Ethernet ports and 1–2 for the Gigabit Ethernet ports. Table 15-1 shows the switch and module combinations and the interface numbers.

Table 15-1 Switch Interface Numbers

Switch Model	Interface Numbering Scheme
IE-2000-4TS-L switch	Fast Ethernet1/1, Fast Ethernet1/2, Fast Ethernet1/3, Fast Ethernet1/4, Gigabit Ethernet1/1, and Gigabit Ethernet1/2
IE-2000-4TS-B switch	Fast Ethernet1/1, Fast Ethernet1/2, Fast Ethernet1/3, Fast Ethernet1/4, Gigabit Ethernet1/1, and Gigabit Ethernet1/2
IE-2000-4T-L switch	Fast Ethernet1/1, Fast Ethernet1/2, Fast Ethernet1/3, Fast Ethernet1/4, Gigabit Ethernet1/1, and Gigabit Ethernet1/2
IE-2000-4T-B switch	Fast Ethernet1/1, Fast Ethernet1/2, Fast Ethernet1/3, Fast Ethernet1/4, Gigabit Ethernet1/1, and Gigabit Ethernet1/2
IE-2000-4TS-G--L switch	Fast Ethernet1/1, Fast Ethernet1/2, Fast Ethernet1/3, Fast Ethernet1/4, Gigabit Ethernet1/1, and Gigabit Ethernet1/2
IE-2000-4TS-G-B switch	Fast Ethernet1/1, Fast Ethernet1/2, Fast Ethernet1/3, Fast Ethernet1/4, Gigabit Ethernet1/1, and Gigabit Ethernet1/2
IE-2000-8TC-L switch	Fast Ethernet1/1, Fast Ethernet1/2, Fast Ethernet1/3, Fast Ethernet1/4, Fast Ethernet1/5, Fast Ethernet1/6, Fast Ethernet1/7, Fast Ethernet1/8, Gigabit Ethernet1/1, and Gigabit Ethernet1/2
IE-2000-8TC-B switch	Fast Ethernet1/1, Fast Ethernet1/2, Fast Ethernet1/3, Fast Ethernet1/4, Fast Ethernet1/5, Fast Ethernet1/6, Fast Ethernet1/7, Fast Ethernet1/8, Gigabit Ethernet1/1, and Gigabit Ethernet1/2
	Fast Ethernet2/1, Fast Ethernet2/2, Fast Ethernet2/3, Fast Ethernet2/4, Fast Ethernet2/5, Fast Ethernet2/6, Fast Ethernet2/7, and Fast Ethernet2/8
	Fast Ethernet3/1, Fast Ethernet3/2, Fast Ethernet3/3, Fast Ethernet3/4, Fast Ethernet3/5, Fast Ethernet3/6, Fast Ethernet3/7, and Fast Ethernet3/8

You can identify physical interfaces by looking at the switch. You can also use the **show** privileged EXEC commands to display information about a specific interface or all the interfaces.

Default Ethernet Interface Settings

For more details on the VLAN parameters listed in the table, see Chapter 17, "Configuring VLANs." For details on controlling traffic to the port, see Chapter 29, "Configuring Port-Based Traffic Control."

Note To configure Layer 2 parameters, if the interface is in Layer 3 mode, you must enter the **switchport** interface configuration command without any parameters to put the interface into Layer 2 mode. This shuts down the interface and then reenables it, which might generate messages on the device to which the interface is connected. When you put an interface that is in Layer 3 mode into Layer 2 mode, the previous configuration information related to the affected interface might be lost, and the interface is returned to its default configuration.

Table 15-2 Default Layer 2 Ethernet Interface Settings

Feature	Default Setting
Operating mode	Layer 2 or switching mode (**switchport** command).
Allowed VLAN range	VLANs 1 to 4096.
Default VLAN (for access ports)	VLAN 1 (Layer 2 interfaces only).
Native VLAN (for IEEE 802.1Q trunks)	VLAN 1 (Layer 2 interfaces only).
VLAN trunking	Switch port mode dynamic auto (supports DTP) (Layer 2 interfaces only).
Port enable state	All ports are enabled.
Port description	None defined.
Speed	Autonegotiate.
Duplex mode	Autonegotiate.
Flow control	Flow control is set to **receive**: **off**. It is always off for sent packets.
EtherChannel (PAgP)	Disabled on all Ethernet ports. Chapter 40, "Configuring EtherChannels."
Port blocking (unknown multicast and unknown unicast traffic)	Disabled (not blocked) (Layer 2 interfaces only).
Broadcast, multicast, and unicast storm control	Disabled.
Protected port	Disabled (Layer 2 interfaces only).
Port security	Disabled (Layer 2 interfaces only).
Port Fast	Disabled.
Auto-MDIX	Enabled. **Note** The switch might not support a prestandard powered device—such as Cisco IP phones and access points that do not fully support IEEE 802.3af—if that powered device is connected to the switch through a crossover cable. This is regardless of whether auto-MIDX is enabled on the switch port.
Keepalive messages	Disabled on SFP module ports; enabled on all other ports.

Interface Speed and Duplex Mode

Depending on the supported port types, Ethernet interfaces on the switch operate at 10, 100, or 1000 Mb/s, or in either full- or half-duplex mode. In full-duplex mode, two stations can send and receive traffic at the same time. Normally, 10-Mb/s ports operate in half-duplex mode, which means that stations can either receive or send traffic.

Switch models can include combinations of Fast Ethernet (10/100-Mb/s) ports, Gigabit Ethernet (10/100/1000-Mb/s) ports, and small form-factor pluggable (SFP) module slots supporting SFP modules.

Speed and Duplex Configuration Guidelines

When configuring an interface speed and duplex mode, note these guidelines:

- Fast Ethernet (10/100-Mb/s) ports support all speed and duplex options.
- Gigabit Ethernet (10/100/1000-Mb/s) ports support all speed options and all duplex options (auto, half, and full). However, Gigabit Ethernet ports operating at 1000 Mb/s do not support half-duplex mode.
- For SFP module ports, the speed and duplex CLI options change depending on the SFP module type:
  - The 1000BASE-*x* (where -*x* is -BX, -CWDM, -LX, -SX, and -ZX) SFP module ports support the **nonegotiate** keyword in the **speed** interface configuration command. Duplex options are not supported.
  - The 1000BASE-T SFP module ports support the same speed and duplex options as the 10/100/1000-Mb/s ports.
  - The 100BASE-*x* (where -*x* is -BX, -CWDM, -LX, -SX, and -ZX) SFP module ports support only 100 Mb/s. These modules support full- and half- duplex options but do not support autonegotiation.

  For information about which SFP modules are supported on your switch, see the product release notes.
- If both ends of the line support autonegotiation, we highly recommend the default setting of **auto** negotiation.
- If one interface supports autonegotiation and the other end does not, configure duplex and speed on both interfaces; do not use the **auto** setting on the supported side.
- When STP is enabled and a port is reconfigured, the switch can take up to 30 seconds to check for loops. The port LED is amber while STP reconfigures.

Caution Changing the interface speed and duplex mode configuration might shut down and reenable the interface during the reconfiguration.

IEEE 802.3x Flow Control

Flow control enables connected Ethernet ports to control traffic rates during congestion by allowing congested nodes to pause link operation at the other end. If one port experiences congestion and cannot receive any more traffic, it notifies the other port by sending a pause frame to stop sending until the condition clears. Upon receipt of a pause frame, the sending device stops sending any data packets, which prevents any loss of data packets during the congestion period.

Note Ports on the switch can receive, but not send, pause frames.

You use the **flowcontrol** interface configuration command to set the interface's ability to **receive** pause frames to **on**, **off**, or **desired**. The default state is **off**.

When set to **desired**, an interface can operate with an attached device that is required to send flow-control packets or with an attached device that is not required to but can send flow-control packets.

These rules apply to flow control settings on the device:

- **receive on** (or **desired**)—The port cannot send pause frames but can operate with an attached device that is required to or can send pause frames; the port can receive pause frames.
- **receive off**—Flow control does not operate in either direction. In case of congestion, no indication is given to the link partner, and no pause frames are sent or received by either device.

Auto-MDIX on an Interface

When automatic medium-dependent interface crossover (auto-MDIX) is enabled on an interface, the interface automatically detects the required cable connection type (straight through or crossover) and configures the connection appropriately. When connecting switches without the auto-MDIX feature, you must use straight-through cables to connect to devices such as servers, workstations, or routers and crossover cables to connect to other switches or repeaters. With auto-MDIX enabled, you can use either type of cable to connect to other devices, and the interface automatically corrects for any incorrect cabling. For more information about cabling requirements, see the hardware installation guide.

Auto-MDIX is enabled by default. When you enable auto-MDIX, you must also set the interface speed and duplex to **auto** so that the feature operates correctly.

Auto-MDIX is supported on all 10/100 and 10/100/1000-Mb/s interfaces. It is not supported on 1000BASE-SX or -LX SFP module interfaces.

SVI Autostate Exclude

Configuring SVI autostate exclude on an access or trunk port in an SVI excludes that port in the calculation of the status of the SVI (up or down line state) even if it belongs to the same VLAN. When the excluded port is in the up state, and all other ports in the VLAN are in the down state, the SVI state is changed to down.

At least one port in the VLAN should be up and not excluded to keep the SVI line state up. You can use this command to exclude the monitoring port status when determining the status of the SVI.

System MTU

The default maximum transmission unit (MTU) size for frames received and transmitted on all interfaces is 1500 bytes. You can increase the MTU size for all interfaces operating at 10 or 100 Mb/s by using the **system mtu** global configuration command. You can increase the MTU size to support jumbo frames on all Gigabit Ethernet interfaces by using the **system mtu jumbo** global configuration command.

You can change the MTU size for routed ports by using the **system mtu routing** global configuration command.

Note You cannot configure a routing MTU size that exceeds the system MTU size. If you change the system MTU size to a value smaller than the currently configured routing MTU size, the configuration change is accepted, but not applied until the next switch reset. When the configuration change takes effect, the routing MTU size automatically defaults to the new system MTU size.

Gigabit Ethernet ports are not affected by the **system mtu** command; 10/100 ports are not affected by the **system mtu jumbo** command. If you d o not configure the **system mtu jumbo** command, the setting of the **system mtu** command applies to all Gigabit Ethernet interfaces.

You cannot set the MTU size for an individual interface; you set it for all 10/100 or all Gigabit Ethernet interfaces. When you change the system or jumbo MTU size, you must reset the switch before the new configuration takes effect.The **system mtu routing** command does not require a switch reset to take effect.

Frames sizes that can be received by the switch CPU are limited to 1998 bytes, no matter what value was entered with the **system mtu** or **system mtu jumbo** commands. Although frames that are forwarded or routed are typically not received by the CPU, in some cases packets are sent to the CPU, such as traffic sent to control traffic, SNMP, Telnet, or routing protocols.

Routed packets are subjected to MTU checks on the output ports. The MTU value used for routed ports is derived from the applied **system mtu** value (not the **system mtu jumbo** value). That is, the routed MTU is never greater than the system MTU for any VLAN. The routing protocols use the system MTU value when negotiating adjacencies and the MTU of the link. For example, the Open Shortest Path First (OSPF) protocol uses this MTU value before setting up an adjacency with a peer router. To view the MTU value for routed packets for a specific VLAN, use the **show platform port-asic mvid** privileged EXEC command.

Note If Layer 2 Gigabit Ethernet interfaces are configured to accept frames greater than the 10/100 interfaces, jumbo frames received on a Layer 2 Gigabit Ethernet interface and sent on a Layer 2 10/100 interface are dropped.

How to Configure Interface Characteristics

Configuring Layer 3 Interfaces

The switch does not support routing, but you can configure an IP address on an interface by using the **no switchport** interface configuration command to facilitate device management.

The switch supports these types of Layer 3 interfaces:

- SVIs: You should configure SVIs for any VLANs for which you want to route traffic. SVIs are created when you enter a VLAN ID following the **interface vlan** global configuration command. To delete an SVI, use the **no interface vlan** global configuration command. You cannot delete interface VLAN 1.

  Note When you create an SVI, it does not become active until it is associated with a physical port. For information about assigning Layer 2 ports to VLANs, see Chapter 17, “Configuring VLANs.”

When configuring SVIs, you can also configure SVI autostate exclude on a port in the SVI to exclude that port from being included in determining SVI line-state status. See the "Configuring SVI Autostate Exclude" section on page 15-17.

- Routed ports: Routed ports are physical ports configured to be in Layer 3 mode by using the **no switchport** interface configuration command.
- Layer 3 EtherChannel ports: EtherChannel interfaces made up of routed ports.

  EtherChannel port interfaces are described in Chapter 40, "Configuring EtherChannels."

A Layer 3 switch can have an IP address assigned to each routed port and SVI.

There is no defined limit to the number of SVIs and routed ports that can be configured in a switch. However, the interrelationship between the number of SVIs and routed ports and the number of other features being configured might have an impact on CPU usage because of hardware limitations. If the switch is using maximum hardware resources, attempts to create a routed port or SVI have these results:

- If you try to create a new routed port, the switch generates a message that there are not enough resources to convert the interface to a routed port, and the interface remains as a switchport.
- If you try to create an extended-range VLAN, an error message is generated, and the extended-range VLAN is rejected.
- If the switch is notified by VLAN Trunking Protocol (VTP) of a new VLAN, it sends a message that there are not enough hardware resources available and shuts down the VLAN. The output of the **show vlan** user EXEC command shows the VLAN in a suspended state.
- If the switch attempts to boot up with a configuration that has more VLANs and routed ports than hardware can support, the VLANs are created, but the routed ports are shut down, and the switch sends a message that this was due to insufficient hardware resources.

All Layer 3 interfaces require an IP address to route traffic. This procedure shows how to configure an interface as a Layer 3 interface and how to assign an IP address to an interface.

Note If the physical port is in Layer 2 mode (the default), you must enter the **no switchport** interface configuration command to put the interface into Layer 3 mode. Entering a **no switchport** command disables and then re-enables the interface, which might generate messages on the device to which the interface is connected. Furthermore, when you put an interface that is in Layer 2 mode into Layer 3 mode, the previous configuration information related to the affected interface might be lost, and the interface is returned to its default configuration.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** {{**fastethernet** | **gigabitethernet**} *interface-id*} | {**vlan** *vlan-id*} | {**port-channel** *port-channel-number*}	Specifies the interface to be configured as a Layer 3 interface, and enters interface configuration mode.
Step 3	**no switchport**	For physical ports only, enter Layer 3 mode.
Step 4	**ip address** *ip_address subnet_mask*	Configures the IP address and IP subnet.
Step 5	**no shutdown**	Enables the interface.
Step 6	**end**	Returns to privileged EXEC mode.

Configuring Interfaces

These general instructions apply to all interface configuration processes.

Step 1 Enter the **configure terminal** command at the privileged EXEC prompt:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)#
```

Step 2 Enter the **interface** global configuration command.

Identify the interface type and the interface number, Gigabit Ethernet port 1 in this example:

```
Switch(config)# interface gigabitethernet1/1
Switch(config-if)#
```

Note Entering a space between the interface type and interface number is optional

Step 3 Follow each **interface** command with the configuration commands that the interface requires. The commands that you enter define the protocols and applications that will run on the interface. The commands are collected and applied to the interface when you enter another interface command or enter **end** to return to privileged EXEC mode.

You can also configure a range of interfaces by using the **interface range** or **interface range macro** global configuration commands. Interfaces configured in a range must be the same type and must be configured with the same feature options.

Step 4 After you configure an interface, verify its status by using the **show** privileged EXEC commands listed in the "Monitoring and Maintaining Interface Characteristics" section on page 15-18.

Enter the **show interfaces** privileged EXEC command to see a list of all interfaces on or configured for the switch. A report is provided for each interface that the device supports or for the specified interface.

Configuring a Range of Interfaces

You can use the **interface range** global configuration command to configure multiple interfaces with the same configuration parameters. When you enter the interface-range configuration mode, all command parameters that you enter are attributed to all interfaces within that range until you exit this mode.

Interface Range Restrictions

- When you use the **interface range** command with port channels, the first and last port-channel number must be active port channels.
- The **interface range** command only works with VLAN interfaces that have been configured with the **interface vlan** command. The **show running-config** privileged EXEC command displays the configured VLAN interfaces. VLAN interfaces not displayed by the **show running-config** command cannot be used with the **interface range** command.

- All interfaces defined as in a range must be the same type (all Fast Ethernet ports, all Gigabit Ethernet ports, all EtherChannel ports, or all VLANs), but you can combine multiple interface types in a macro.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface range** {*port-range* | **macro** *macro_name*}	Specifies the range of interfaces (VLANs or physical ports) to be configured, and enters interface-range configuration mode. • **interface range**—Configures up to five port ranges or a previously defined macro. • **macro** *macro_name*—Specifies the 32-character maximum character string. • In a comma-separated *port-range*, you must enter the interface type for each entry and enter spaces before and after the comma. • In a hyphen-separated *port-range*, you do not need to reenter the interface type, but you must enter a space before the hyphen.
Step 3		Uses the normal configuration commands to apply the configuration parameters to all interfaces in the range. Each command is executed as it is entered.
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show interfaces** [*interface-id*]	Verifies the configuration of the interfaces in the range.
Step 6	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Configuring and Using Interface Range Macros

Before You Begin

You can create an interface range macro to automatically select a range of interfaces for configuration. Before you can use the **macro** keyword in the **interface range macro** global configuration command string, you must use the **define interface-range** global configuration command to define the macro.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**define interface-range** *macro_name interface-range*	Defines the interface-range macro, and saves it in NVRAM. • **macro** *macro_name*—Specifies the 32-character maximum character string. • A macro can contain up to five comma-separated interface ranges. • *interface-range*—Consists of the same port type.
Step 3	**interface range macro** *macro_name*	Selects the interface range to be configured using the values saved in the interface-range macro called *macro_name.* You can now use the normal configuration commands to apply the configuration to all interfaces in the defined macro.

	Command	Purpose
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show running-config | include define**	Shows the defined interface range macro configuration.

Configuring Ethernet Interfaces

Setting the Type of a Dual-Purpose Uplink Port

Perform this task to select which dual-purpose uplink to activate so that you can set the speed and duplex. This procedure is optional.

	Command	Purpose
Step 1	**configure terminal**	Enters globals configuration mode.
Step 2	**interface** *interface-id*	Specifies the dual-purpose uplink port to be configured, and enters interface configuration mode.
Step 3	**media-type** {**auto-select** | **rj45** | **sfp**}	Selects the interface and type of a dual-purpose uplink port. The keywords have these meanings: • **auto-select**—The switch dynamically selects the type. When link up is achieved, the switch disables the other type until the active link goes down. When the active link goes down, the switch enables both types until one of them links up. In auto-select mode, the switch configures both types with autonegotiation of speed and duplex (the default). Depending on the type of installed SFP module, the switch might not be able to dynamically select it. • **rj45**—The switch disables the SFP module interface. If you connect an SFP module to this port, it cannot attain a link even if the RJ-45 side is down or is not connected. In this mode, the dual-purpose port behaves like a 10/100/1000BASE-TX interface. You can configure the speed and duplex settings consistent with this interface type. • **sfp**—The switch disables the RJ-45 interface. If you connect a cable to the RJ-45 port, it cannot attain a link even if the SFP module side is down or if the SFP module is not present. Based on the type of installed SFP module, you can configure the speed and duplex settings consistent with this interface type.
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show interfaces** *interface-id* **transceiver properties**	Verifies your setting.

Setting the Interface Speed and Duplex Parameters

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the physical interface to be configured, and enters interface configuration mode.
Step 3	**speed** {**10** | **100** | **1000** | **auto** [**10** | **100** | **1000**] | **nonegotiate**}	Enters the appropriate speed parameter for the interface: • **10**, **100**, or **1000**—Sets a specific speed for the interface. The **1000** keyword is available only for 10/100/1000 Mb/s ports. • **auto**—Enables the interface to autonegotiate speed with the connected device. If you use the **10**, **100**, or the **1000** keywords with the **auto** keyword, the port autonegotiates only at the specified speeds. • **nonegotiate**—Available only for SFP module ports. SFP module ports operate only at 1000 Mb/s but can be configured to not negotiate if connected to a device that does not support autonegotiation.
Step 4	**duplex** {**auto** | **full** | **half**}	Enters the duplex parameter for the interface. Enables half-duplex mode (for interfaces operating only at 10 or 100 Mb/s). You cannot configure half-duplex mode for interfaces operating at 1000 Mb/s.
Step 5	**end**	Returns to privileged EXEC mode.
Step 6	**show interfaces** *interface-id*	Displays the interface speed and duplex mode configuration.

Configuring IEEE 802.3x Flow Control

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the physical interface to be configured, and enter interface configuration mode.
Step 3	**flowcontrol** {**receive**} {**on** | **off** | **desired**}	Configures the flow control mode for the port.
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show interfaces** *interface-id*	Verifies the interface flow control settings.

Configuring Auto-MDIX on an Interface

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the physical interface to be configured, and enters interface configuration mode.
Step 3	**speed auto**	Configures the interface to autonegotiate speed with the connected device.
Step 4	**duplex auto**	Configures the interface to autonegotiate duplex mode with the connected device.
Step 5	**mdix auto**	Enables auto-MDIX on the interface.
Step 6	**end**	Returns to privileged EXEC mode.
Step 7	**show controllers ethernet-controller** *interface-id* **phy**	Verifies the operational state of the auto-MDIX feature on the interface.

Adding a Description for an Interface

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface for which you are adding a description, and enters interface configuration mode.
Step 3	**description** *string*	Adds a description (up to 240 characters) for an interface.
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show interfaces** *interface-id* **description** or **show running-config**	Verifies your entry.

Configuring SVI Autostate Exclude

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies a Layer 2 interface (physical port or port channel), and enters interface configuration mode.
Step 3	**switchport autostate exclude**	Excludes the access or trunk port when defining the status of an SVI line state (up or down)
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show running config interface** *interface-id* **show interface** *interface-id* **switchport**	(Optional) Shows the running configuration. Verifies the configuration.

Configuring the System MTU

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**system mtu** *bytes*	(Optional) Changes the MTU size for all interfaces on the switch that are operating at 10 or 100 Mb/s. The range is 1500 to 1998 bytes; the default is 1500 bytes.
Step 3	**system mtu jumbo** *bytes*	(Optional) Changes the MTU size for all Gigabit Ethernet interfaces on the switch. The range is 1500 to 9000 bytes; the default is 1500 bytes.
Step 4	**system mtu routing** *bytes*	(Optional) Changes the system MTU for routed ports. The range is 1500 to the system MTU value, the maximum MTU that can be routed for all ports. Although larger packets can be accepted, they cannot be routed.
Step 5	**end**	Returns to privileged EXEC mode.
Step 6	**copy running-config startup-config**	Saves your entries in the configuration file.
Step 7	**reload**	Reloads the operating system.
Step 8	**show system mtu**	(Optional) Verifies your settings.

Monitoring and Maintaining Interface Characteristics

Monitoring Interface Status

Table 15-3 Show Commands for Interfaces

Command	Purpose
show interfaces [*interface-id*]	(Optional) Displays the status and configuration of all interfaces or a specific interface. **Note** A disabled interface is shown as *administratively down* in the display.
show interfaces *interface-id* **status** [**err-disabled**]	(Optional) Displays interface status or a list of interfaces in an error-disabled state.
show interfaces [*interface-id*] **switchport**	(Optional) Displays administrative and operational status of switching ports. You can use this command to find out if a port is in routing or in switching mode.
show interfaces [*interface-id*] **description**	(Optional) Displays the description configured on an interface or all interfaces and the interface status.

Table 15-3 Show Commands for Interfaces (continued)

Command	Purpose
show ip interface [*interface-id*]	(Optional) Displays the usability status of all interfaces configured for IP routing or the specified interface.
show interface [*interface-id*] **stats**	(Optional) Displays the input and output packets by the switching path for the interface.
show interfaces transceiver properties	(Optional) Displays speed and duplex settings on the interface.
show interfaces transceiver detail	(Optional) Displays temperature, voltage, or amount of current on the interface.
show interfaces [*interface-id*] [{**transceiver properties** | **detail**}] *module number*]	Displays physical and operational status about an SFP module.
show running-config interface [*interface-id*]	Displays the running configuration in RAM for the interface.
show version	Displays the hardware configuration, software version, the names and sources of configuration files, and the boot images.
show controllers ethernet-controller *interface-id* **phy**	Displays the operational state of the auto-MDIX feature on the interface.

Clearing and Resetting Interfaces and Counters

Table 15-4 Clear Commands for Interfaces

Command	Purpose
clear counters [*interface-id*]	Clears interface counters. **Note** This command does not clear counters retrieved by using Simple Network Management Protocol (SNMP), but only those seen with the **show interface** privileged EXEC command.
clear interface *interface-id*	Resets the hardware logic on an interface.
clear line [*number* | **console 0** | **vty** *number*]	Resets the hardware logic on an asynchronous serial line.

Shutting Down and Restarting the Interface

Shutting down an interface disables all functions on the specified interface and marks the interface as unavailable on all monitoring command displays. This information is communicated to other network servers through all dynamic routing protocols. The interface is not mentioned in any routing updates.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** {**vlan** *vlan-id*} | {{**fastethernet** | **gigabitethernet**} *interface-id*} | {**port-channel** *port-channel-number*}	Selects the interface to be configured.

	Command	Purpose
Step 3	**shutdown**	Shuts down an interface. **Note** Use the **no shutdown** interface configuration command to restart the interface.
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show running-config**	Verifies your entry.

Configuration Examples for Configuring Interface Characteristics

Configuring the Interface Range: Examples

This example shows how to use the **interface range** global configuration command to set the speed on ports 1 to 2 to 100 Mb/s:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet1/1 - 2
Switch(config-if-range)# speed 100
```

This example shows how to use a comma to add different interface type strings to the range to enable Fast Ethernet ports 1 to 3 and Gigabit Ethernet ports 1 and 2 to receive flow-control pause frames:

```
Switch# configure terminal
Switch(config)# interface range fastethernet1/1 - 3, gigabitethernet1/1 - 2
Switch(config-if-range)# flowcontrol receive on
```

If you enter multiple configuration commands while you are in interface-range mode, each command is executed as it is entered. The commands are not batched and executed after you exit interface-range mode. If you exit interface-range configuration mode while the commands are being executed, some commands might not be executed on all interfaces in the range. Wait until the command prompt reappears before exiting interface-range configuration mode.

Configuring Interface Range Macros: Examples

This example shows how to define an interface-range named *enet_list* to include ports 1 and 2 and to verify the macro configuration:

```
Switch# configure terminal
Switch(config)# define interface-range enet_list gigabitethernet1/1 - 2
Switch(config)# end
Switch# show running-config | include define
Switch# define interface-range enet_list gigabitethernet1/1 - 2
```

This example shows how to create a multiple-interface macro named *macro1*:

```
Switch# configure terminal
Switch(config)# define interface-range macro1 fastethernet1/1 - 2, gigabitethernet1/1 - 2
Switch(config)# end
```

This example shows how to enter interface-range configuration mode for the interface-range macro *enet_list*:

```
Switch# configure terminal
Switch(config)# interface range macro enet_list
Switch(config-if-range)#
```

This example shows how to delete the interface-range macro *enet_list* and to verify that it was deleted.

```
Switch# configure terminal
Switch(config)# no define interface-range enet_list
Switch(config)# end
Switch# show run | include define
Switch#
```

Setting Speed and Duplex Parameters: Example

This example shows how to set the interface speed to 10 Mb/s and the duplex mode to half on a 10/100 Mb/s port:

```
Switch# configure terminal
Switch(config)# interface fasttethernet1/3
Switch(config-if)# speed 10
Switch(config-if)# duplex half
```

This example shows how to set the interface speed to 100 Mb/s on a 10/100/1000 Mb/s port:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/2
Switch(config-if)# speed 100
```

Enabling auto-MDIX: Example

This example shows how to enable auto-MDIX on a port:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# speed auto
Switch(config-if)# duplex auto
Switch(config-if)# mdix auto
Switch(config-if)# end
```

Adding a Description on a Port: Example

This example shows how to add a description on a port and how to verify the description:

```
Switch# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet1/2
Switch(config-if)# description Connects to Marketing
Switch(config-if)# end
Switch# show interfaces gigabitethernet1/2 description
Interface Status         Protocol Description
Gi1/2   admin down     down     Connects to Marketing
```

Configuring SVI Autostate Exclude: Example

This example shows how to configure an access or trunk port in an SVI to be excluded from the status calculation:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet1/2
Switch(config-if)# switchport autostate exclude
Switch(config-if)# exit
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
Cisco IOS interface commands	*Cisco IOS Interface Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

CHAPTER 16

Configuring Smartports Macros

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Information About Configuring Smartports Macros

Smartports macros provide a convenient way to save and share common configurations. You can use Smartports macros to enable features and settings based on the location of a switch in the network and for mass configuration deployments across the network.

Each Smartports macro is a set of CLI commands that you define. Smartports macros do not contain new CLI commands; they are simply a group of existing CLI commands.

When you apply a Smartports macro to an interface, the CLI commands within the macro are configured on the interface. When the macro is applied to an interface, the existing interface configurations are not lost. The new commands are added to the interface and are saved in the running configuration file.

How to Configure Smartports Macros

Default Smartports Settings

There are no Smartports macros enabled on the switch.

Table 16-1 Default Smartports Macros

Macro Name[1]	Description
cisco-ie-global	Use this global configuration macro to configure the switch settings for the industrial Ethernet environment. This macro is automatically applied when you use Express Setup to initially configure the switch. **Note** You must first apply the **cisco-ie-global** macro for the **cisco-ethernetip** macro to work properly.
cisco-desktop	Use this interface configuration macro for increased network security and reliability when connecting a desktop device, such as a PC, to a switch port. This macro is optimized for industrial automation traffic.
cisco-phone	Use this interface configuration macro when connecting a desktop device such as a PC with a Cisco IP phone to a switch port. This macro is an extension of the **cisco-ie-desktop** macro and provides the same security and resiliency features, but with the addition of dedicated voice VLANs to ensure proper treatment of delay-sensitive voice traffic. This macro is optimized for industrial automation traffic.
cisco-ie-switch	Use this interface configuration macro when connecting an access switch and a distribution switch or between access switches connected using small form-factor pluggable (SFP) modules. This macro is optimized for industrial automation traffic.
cisco-router	Use this interface configuration macro when connecting the switch and a WAN router. This macro is optimized for industrial automation traffic.
cisco-ethernetip	Use this interface configuration macro when connecting the switch to an EtherNet IP device. **Note** You must first apply the **cisco-ie-global** macro for the **cisco-ethernetip** macro to work properly.
cisco-ie-qos-map-setup	Use this global configuration macro to configure the QoS policy map for for the industrial Ethernet environment.
cisco-ie-qos-queue-setup	Use this global configuration macro to configure the QoS policy map for for the industrial Ethernet environment.

1. Cisco-default Smartports macros vary, depending on the software version running on your switch.

Smartports Configuration Guidelines

- When a macro is applied globally to a switch or to a switch interface, all of the existing configurations on the interface are retained. This is helpful when applying an incremental configuration.
- If a command fails because of a syntax or a configuration error, the macro continues to apply the remaining commands. You can use the **macro global trace** *macro-name* global configuration command or the **macro trace** *macro-name* interface configuration command to apply and debug a macro to find any syntax or configuration errors.
- Some CLI commands are specific to certain interface types. If you apply a macro to an interface that does not accept the configuration, the macro fails the syntax or the configuration check, and the switch returns an error message.
- Applying a macro to an interface range is the same as applying a macro to a single interface. When you use an interface range, the macro is applied sequentially to each interface within the range. If a macro command fails on one interface, it is still applied to the remaining interfaces.

- When you apply a macro to a switch or a switch interface, the macro name is automatically added to the switch or interface. You can display the applied commands and macro names by using the **show running-config** user EXEC command.

Applying Smartports Macros

	Command	Purpose
Step 1	**show parser macro**	Displays the Cisco-default Smartports macros embedded in the switch software.
Step 2	**show parser macro name** *macro-name*	Displays the specific macro that you want to apply.
Step 3	**configure terminal**	Enters global configuration mode.
Step 4	**macro global** {**apply** | **trace**} *macro-name* [**parameter** {*value*}] [**parameter** {*value*}] [**parameter** {*value*}]	Applies each individual command defined in the macro to the switch by entering **macro global apply** *macro-name*. Specifies **macro global trace** *macro-name* to apply and to debug a macro to find any syntax or configuration errors. Appends the macro with the required values by using the **parameter** *value* keywords. Keywords that begin with **$** require a unique parameter value. You can use the **macro global apply** *macro-name* **?** command to display a list of any required values for the macro. If you apply a macro without entering the keyword values, the commands are invalid and are not applied. (Optional) Specifies unique parameter values that are specific to the switch. You can enter up to three keyword-value pairs. Parameter keyword matching is case sensitive. The corresponding value replaces all matching occurrences of the keyword.
Step 5	**interface** *interface-id*	(Optional) Enters interface configuration mode and specifies the interface on which to apply the macro.
Step 6	**default interface** *interface-id*	(Optional) Clears all configuration from the specified interface.
Step 7	**macro** {**apply** | **trace**} *macro-name* [**parameter** {*value*}] [**parameter** {*value*}] [**parameter** {*value*}]	Applies each individual command defined in the macro to the port by entering **macro global apply** *macro-name*. Specifies **macro global trace** *macro-name* to apply and to debug a macro to find any syntax or configuration errors. Appends the macro with the required values by using the **parameter** *value* keywords. Keywords that begin with **$** require a unique parameter value. You can use the **macro global apply** *macro-name* **?** command to display a list of any required values for the macro. If you apply a macro without entering the keyword values, the commands are invalid and are not applied. (Optional) Specifies unique parameter values that are specific to the switch. You can enter up to three keyword-value pairs. Parameter keyword matching is case sensitive. The corresponding value replaces all matching occurrences of the keyword.
Step 8	**end**	Returns to privileged EXEC mode.

	Command	Purpose
Step 9	**show running-config interface** *interface-id*	Verifies that the macro is applied to an interface.
Step 10	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Monitoring and Maintaining Smartports Macros

Table 16-2 Commands for Displaying Smartports Macros

Command	Purpose
show parser macro	Displays all Smartports macros.
show parser macro name *macro-name*	Displays a specific Smartports macro.
show parser macro brief	Displays the Smartports macro names.
show parser macro description [**interface** *interface-id*]	Displays the Smartports macro description for all interfaces or for a specified interface.

Configuration Examples for Smartports Macros

Applying the Smartports Macro: Examples

This example shows how to display the **cisco-ie-desktop** macro, how to apply the macro and to set the access VLAN ID to 25 on an interface:

```
Switch# show parser macro name cisco-ie-desktop
-------------------------------------------------------------
Macro name : cisco-ie-desktop
Macro type : default interface
# macro keywords ACCESS_VLAN
#macro name cisco-ie-desktop
switchport mode access
switchport access vlan ACCESS_VLAN
switchport port-security
switchport port-security maximum 1
switchport port-security aging time 2
switchport port-security violation restrict
switchport port-security aging type inactivity
spanning-tree portfast
spanning-tree bpduguard enable
no macro description
macro description cisco-ie-desktop
-------------------------------------------------------------
Switch#
Switch# configure terminal
Switch(config)# interface gigabitethernet1/4
Switch(config-if)# macro apply cisco-ie-desktop $AVID 25
```

This example shows how to display the **cisco-ethernetip** macro and how to apply it to an interface:

```
Switch# show parser macro name cisco-ethernetip
Macro name : cisco-ie-global
```

```
Macro type : default interface
#macro name cisco-ethernetip
#macro keywords ACCESS_VLAN
#macro description cisco-ethernetip
switchport host
switchport access vlan ACCESS-VLAN
storm-control broadcast level 3.00 1.00
service-policy input CIP-Traffic
#service-policy input 1588

Switch# configure terminal
Switch(config)# interface fastethernet 1/1
Switch(config-if)# macro apply cisco-ethernetip ACCESS_VLAN 1
switchport mode will be set to access
spanning-tree portfast will be enabled
channel group will be disabled
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 17

Configuring VLANs

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Information About Configuring VLANs

VLANs

A VLAN is a switched network that is logically segmented by function, project team, or application, without regard to the physical locations of the users. VLANs have the same attributes as physical LANs, but you can group end stations even if they are not physically located on the same LAN segment. Any switch port can belong to a VLAN, and unicast, broadcast, and multicast packets are forwarded and flooded only to end stations in the VLAN. Each VLAN is considered a logical network, and packets destined for stations that do not belong to the VLAN must be forwarded through a router or a switch supporting fallback bridging, as shown in Figure 17-1. Because a VLAN is considered a separate logical network, it contains its own bridge Management Information Base (MIB) information and can support its own implementation of spanning tree. See Chapter 20, "Configuring STP."

Note Before you create VLANs, you must decide whether to use VLAN Trunking Protocol (VTP) to maintain global VLAN configuration for your network. For more information on VTP, see Chapter 18, "Configuring VTP."

Figure 17-1 VLANs as Logically Defined Networks

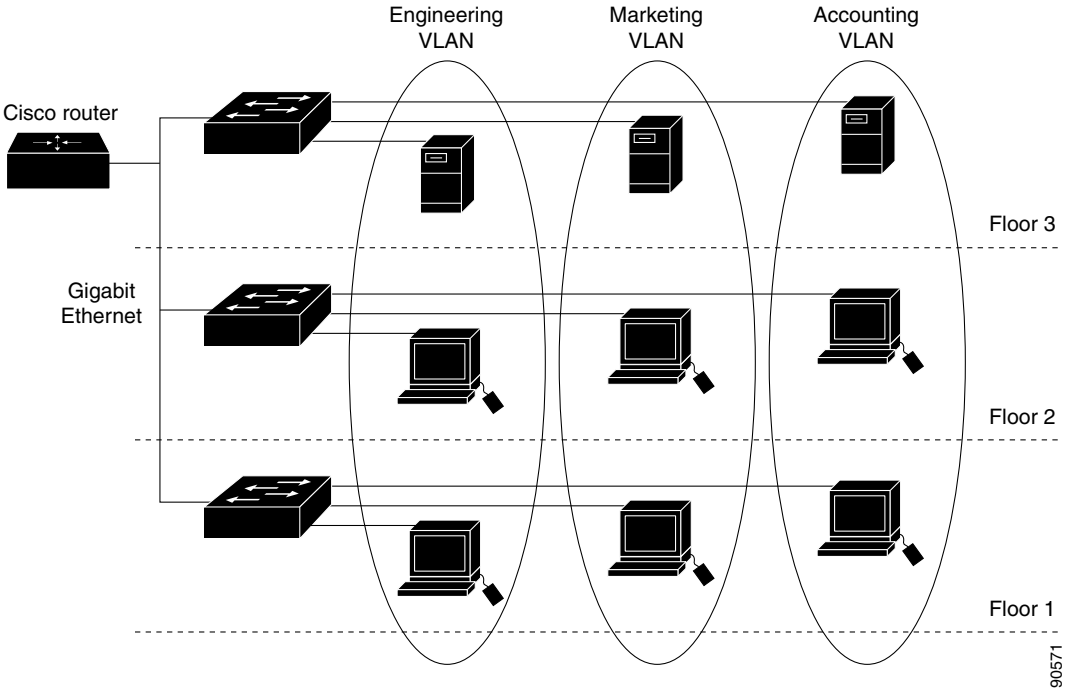

VLANs are often associated with IP subnetworks. For example, all the end stations in a particular IP subnet belong to the same VLAN. Interface VLAN membership on the switch is assigned manually on an interface-by-interface basis. When you assign switch interfaces to VLANs by using this method, it is known as interface-based, or static, VLAN membership.

Traffic between VLANs must be routed or fallback bridged. The switch can route traffic between VLANs by using switch virtual interfaces (SVIs). An SVI must be explicitly configured and assigned an IP address to route traffic between VLANs.

Note If you plan to configure many VLANs on the switch and to not enable routing, you can use the **sdm prefer vlan** global configuration command to set the Switch Database Management (sdm) feature to the VLAN template, which configures system resources to support the maximum number of unicast MAC addresses. For more information on the SDM templates, see Chapter 11, "Configuring SDM Templates," or see the **sdm prefer** command in the command reference for this release.

Supported VLANs

The switch supports VLANs in VTP client, server, and transparent modes. VLANs are identified by a number from 1 to 4096. VLAN IDs 1002 through 1005 are reserved for Token Ring and FDDI VLANs.

VTP version 1 and version 2 support only normal-range VLANs (VLAN IDs 1 to 1005). In these versions, the switch must be in VTP transparent mode when you create VLAN IDs from 1006 to 4096.

This release supports VTP version 3. VTP version 3 supports the entire VLAN range (VLANs 1 to 4096). Extended range VLANs (VLANs 1006 to 4096) are supported only in VTP version 3. You cannot convert from VTP version 3 to VTP version 2 if extended VLANs are configured in the domain.

Although the switch supports a total of 1005 (normal range and extended range) VLANs, the number of routed ports, SVIs, and other configured features affects the use of the switch hardware.

The switch supports per-VLAN spanning-tree plus (PVST+) or rapid PVST+ with a maximum of 128 spanning-tree instances. One spanning-tree instance is allowed per VLAN. See the "Normal-Range VLAN Configuration Guidelines" section on page 17-6 for more information about the number of spanning-tree instances and the number of VLANs.

VLAN Port Membership Modes

You configure a port to belong to a VLAN by assigning a membership mode that specifies the kind of traffic the port carries and the number of VLANs to which it can belong. Table 17-1 lists the membership modes and membership and VTP characteristics.

Table 17-1 Port Membership Modes and Characteristics

Membership Mode	VLAN Membership Characteristics	VTP Characteristics
Static-access	A static-access port can belong to one VLAN and is manually assigned to that VLAN. For more information, see the "Assigning Static-Access Ports to a VLAN" section on page 17-17.	VTP is not required. If you do not want VTP to globally propagate information, set the VTP mode to transparent. To participate in VTP, there must be at least one trunk port on the switch connected to a trunk port of a second switch.
Trunk (ISL or IEEE 802.1Q)	A trunk port is a member of all VLANs by default, including extended-range VLANs, but membership can be limited by configuring the allowed-VLAN list. You can also modify the pruning-eligible list to block flooded traffic to VLANs on trunk ports that are included in the list. For information about configuring trunk ports, see the "Configuring an Ethernet Interface as a Trunk Port" section on page 17-19.	VTP is recommended but not required. VTP maintains VLAN configuration consistency by managing the addition, deletion, and renaming of VLANs on a network-wide basis. VTP exchanges VLAN configuration messages with other switches over trunk links.

Membership Mode	VLAN Membership Characteristics	VTP Characteristics
Dynamic access	A dynamic-access port can belong to one VLAN and is dynamically assigned by a VMPS (VLAN Membership Policy Server). The VMPS can be a Catalyst 5000 or Catalyst 6500 series switch, for example, but never an IE 2000switch. The IE 2000 switch is a VMPS client. You can have dynamic-access ports and trunk ports on the same switch, but you must connect the dynamic-access port to an end station or hub and not to another switch. For configuration information, see the "Configuring Dynamic-Access Ports on VMPS Clients" section on page 17-23.	VTP is required. Configure the VMPS and the client with the same VTP domain name. To participate in VTP, at least one trunk port on the switch must be connected to a trunk port of a second switch.
Voice VLAN	A voice VLAN port is an access port attached to a Cisco IP Phone, configured to use one VLAN for voice traffic and another VLAN for data traffic from a device attached to the phone. For more information about voice VLAN ports, see Chapter 19, "Configuring Voice VLAN."	VTP is not required; it has no effect on a voice VLAN.

For more detailed definitions of access and trunk modes and their functions, see Table 17-3 on page 17-10.

When a port belongs to a VLAN, the switch learns and manages the addresses associated with the port on a per-VLAN basis. For more information, see the "Changing the Address Aging Time" section on page 7-13.

Normal-Range VLANs

Normal-range VLANs are VLANs with VLAN IDs 1 to 1005. If the switch is in VTP server or VTP transparent mode, you can add, modify or remove configurations for VLANs 2 to 1001 in the VLAN database. (VLAN IDs 1 and 1002 to 1005 are automatically created and cannot be removed.)

Configurations for VLAN IDs 1 to 1005 are written to the *vlan.dat* file (VLAN database), and you can display them by entering the **show vlan** privileged EXEC command. The *vlan.dat* file is stored in flash memory.

Caution You can cause inconsistency in the VLAN database if you attempt to manually delete the *vlan.dat* file. If you want to modify the VLAN configuration, use the commands described in these sections and in the command reference for this release. To change the VTP configuration, see Chapter 18, "Configuring VTP."

You use the interface configuration mode to define the port membership mode and to add and remove ports from VLANs. The results of these commands are written to the running-configuration file, and you can display the file by entering the **show running-config** privileged EXEC command.

You can set these parameters when you create a new normal-range VLAN or modify an existing VLAN in the VLAN database:

- VLAN ID
- VLAN name
- VLAN type (Ethernet, Fiber Distributed Data Interface [FDDI], FDDI network entity title [NET], TrBRF, or TrCRF, Token Ring, Token Ring-Net)
- VLAN state (active or suspended)
- Maximum transmission unit (MTU) for the VLAN
- Security Association Identifier (SAID)
- Bridge identification number for TrBRF VLANs
- Ring number for FDDI and TrCRF VLANs
- Parent VLAN number for TrCRF VLANs
- Spanning Tree Protocol (STP) type for TrCRF VLANs
- VLAN number to use when translating from one VLAN type to another

You configure VLANs in **vlan** global configuration command by entering a VLAN ID. Enter a new VLAN ID to create a VLAN, or enter an existing VLAN ID to modify that VLAN. You can use the default VLAN configuration (Table 17-2) or enter multiple commands to configure the VLAN. For more information about commands available in this mode, see the **vlan** global configuration command description in the command reference for this release. When you have finished the configuration, you must exit VLAN configuration mode for the configuration to take effect. To display the VLAN configuration, enter the **show vlan** privileged EXEC command.

The configurations of VLAN IDs 1 to 1005 are always saved in the VLAN database (vlan.dat file). If the VTP mode is transparent, they are also saved in the switch running configuration file. You can enter the **copy running-config startup-config** privileged EXEC command to save the configuration in the startup configuration file. To display the VLAN configuration, enter the **show vlan** privileged EXEC command.

When you save VLAN and VTP information (including extended-range VLAN configuration information) in the startup configuration file and reboot the switch, the switch configuration is selected as follows:

- If the VTP mode is transparent in the startup configuration, and the VLAN database and the VTP domain name from the VLAN database matches that in the startup configuration file, the VLAN database is ignored (cleared), and the VTP and VLAN configurations in the startup configuration file are used. The VLAN database revision number remains unchanged in the VLAN database.
- If the VTP mode or domain name in the startup configuration does not match the VLAN database, the domain name and VTP mode and configuration for the first 1005 VLANs use the VLAN database information.
- In VTP versions 1 and 2, if VTP mode is server, the domain name and VLAN configuration for only the first 1005 VLANs use the VLAN database information. VTP version 3 also supports VLANs 1006 to 4096.

Token Ring VLANs

Although the switch does not support Token Ring connections, a remote device such as a Catalyst 6500 series switch with Token Ring connections could be managed from one of the supported switches. Switches running VTP Version 2 advertise information about these Token Ring VLANs:

- Token Ring TrBRF VLANs

- Token Ring TrCRF VLANs

For more information on configuring Token Ring VLANs, see the *Catalyst 6500 Series Software Configuration Guide*.

Normal-Range VLAN Configuration Guidelines

Follow these guidelines when creating and modifying normal-range VLANs in your network:

- The switch supports 1005 VLANs in VTP client, server, and transparent modes.
- Normal-range VLANs are identified with a number between 1 and 1001. VLAN numbers 1002 through 1005 are reserved for Token Ring and FDDI VLANs.
- VLAN configuration for VLANs 1 to 1005 are always saved in the VLAN database. If the VTP mode is transparent, VTP and VLAN configuration are also saved in the switch running configuration file.
- With VTP versions 1 and 2, the switch supports VLAN IDs 1006 through 4096 only in VTP transparent mode (VTP disabled). These are extended-range VLANs and configuration options are limited. Extended-range VLANs created in VTP transparent mode are not saved in the VLAN database and are not propagated. VTP version 3 supports extended range VLAN (VLANs 1006 to 4096) database propagation. If extended VLANs are configured, you cannot convert from VTP version 3 to version 1 or 2. See the "Creating an Extended-Range VLAN" section on page 17-18.
- Before you can create a VLAN, the switch must be in VTP server mode or VTP transparent mode. If the switch is a VTP server, you must define a VTP domain or VTP will not function.
- The switch does not support Token Ring or FDDI media. The switch does not forward FDDI, FDDI-Net, TrCRF, or TrBRF traffic, but it does propagate the VLAN configuration through VTP.
- The switch supports 128 spanning-tree instances. If a switch has more active VLANs than supported spanning-tree instances, spanning tree can be enabled on 128 VLANs and is disabled on the remaining VLANs. If you have already used all available spanning-tree instances on a switch, adding another VLAN anywhere in the VTP domain creates a VLAN on that switch that is not running spanning-tree. If you have the default allowed list on the trunk ports of that switch (which is to allow all VLANs), the new VLAN is carried on all trunk ports. Depending on the topology of the network, this could create a loop in the new VLAN that would not be broken, particularly if there are several adjacent switches that all have run out of spanning-tree instances. You can prevent this possibility by setting allowed lists on the trunk ports of switches that have used up their allocation of spanning-tree instances.

  If the number of VLANs on the switch exceeds the number of supported spanning-tree instances, we recommend that you configure the IEEE 802.1s Multiple STP (MSTP) on your switch to map multiple VLANs to a single spanning-tree instance. For more information about MSTP, see Chapter 21, "Configuring MSTP."

Default Ethernet VLAN Configuration

Note The switch supports Ethernet interfaces exclusively. Because FDDI and Token Ring VLANs are not locally supported, you only configure FDDI and Token Ring media-specific characteristics for VTP global advertisements to other switches.

Table 17-2 Ethernet VLAN Defaults and Ranges

Parameter	Default	Range
VLAN ID	1	1 to 4096. **Note** Extended-range VLANs (VLAN IDs 1006 to 4096) are only saved in the VLAN database in VTP version 3.
VLAN name	*VLANxxxx*, where *xxxx* represents four numeric digits (including leading zeros) equal to the VLAN ID number	No range
IEEE 802.10 SAID	100001 (100000 plus the VLAN ID)	1 to 4294967294
MTU size	1500	1500 to 18190
Translational bridge 1	0	0 to 1005
Translational bridge 2	0	0 to 1005
VLAN state	active	active, suspend
Remote SPAN	disabled	enabled, disabled

Ethernet VLANs

Each Ethernet VLAN in the VLAN database has a unique, 4-digit ID that can be a number from 1 to 1001. VLAN IDs 1002 to 1005 are reserved for Token Ring and FDDI VLANs. To create a normal-range VLAN to be added to the VLAN database, assign a number and name to the VLAN.

Note With VTP version 1 and 2, if the switch is in VTP transparent mode, you can assign VLAN IDs greater than 1006, but they are not added to the VLAN database. See the "Creating an Extended-Range VLAN" section on page 17-18.

For the list of default parameters that are assigned when you add a VLAN, see the "Normal-Range VLANs" section on page 17-4.

VLAN Removal

When you delete a VLAN from a switch that is in VTP server mode, the VLAN is removed from the VLAN database for all switches in the VTP domain. When you delete a VLAN from a switch that is in VTP transparent mode, the VLAN is deleted only on that specific switch.

You cannot delete the default VLANs for the different media types: Ethernet VLAN 1 and FDDI or Token Ring VLANs 1002 to 1005.

Caution When you delete a VLAN, any ports assigned to that VLAN become inactive. They remain associated with the VLAN (and thus inactive) until you assign them to a new VLAN.

Static-Access Ports for a VLAN

You can assign a static-access port to a VLAN without having VTP globally propagate VLAN configuration information by disabling VTP (VTP transparent mode).

If you are assigning a port on a cluster member switch to a VLAN, first use the **rcommand** privileged EXEC command to log in to the cluster member switch.

Note If you assign an interface to a VLAN that does not exist, the new VLAN is created. (See the "Creating or Modifying an Ethernet VLAN" section on page 17-17.)

Extended-Range VLANs

With VTP version 1 and version 2, when the switch is in VTP transparent mode (VTP disabled), you can create extended-range VLANs (in the range 1006 to 4096). VTP version supports extended-range VLANs in server or transparent move. Extended-range VLANs enable service providers to extend their infrastructure to a greater number of customers. The extended-range VLAN IDs are allowed for any switchport commands that allow VLAN IDs.

With VTP version 1 or 2, extended-range VLAN configurations are not stored in the VLAN database, but because VTP mode is transparent, they are stored in the switch running configuration file, and you can save the configuration in the startup configuration file by using the **copy running-config startup-config** privileged EXEC command. Extended-range VLANs created in VTP version 3 are stored in the VLAN database.

Default VLAN Configuration

See Table 17-2 on page 17-7 for the default configuration for Ethernet VLANs. You can change only the MTU size, private VLAN, and the remote SPAN configuration state on extended-range VLANs; all other characteristics must remain at the default state.

Extended-Range VLAN Configuration Guidelines

Follow these guidelines when creating extended-range VLANs:

- VLAN IDs in the extended range are not saved in the VLAN database and are not recognized by VTP unless the switch is running VTP version 3.
- You cannot include extended-range VLANs in the pruning eligible range.
- In VTP version 1 and 2, a switch must be in VTP transparent mode when you create extended-range VLANs. If VTP mode is server or client, an error message is generated, and the extended-range VLAN is rejected. VTP version 3 supports extended VLANs in server and transparent modes.
- For VTP version 1 or 2, you can set the VTP mode to transparent in global configuration mode. See "Adding a VTP Client Switch to a VTP Domain" section on page 18-10. You should save this configuration to the startup configuration so that the switch boots up in VTP transparent mode. Otherwise, you lose the extended-range VLAN configuration if the switch resets. If you create extended-range VLANs in VTP version 3, you cannot convert to VTP version 1 or 2.
- STP is enabled by default on extended-range VLANs, but you can disable it by using the **no spanning-tree vlan** *vlan-id* global configuration command. When the maximum number of spanning-tree instances are on the switch, spanning tree is disabled on any newly created VLANs. If the number of VLANs on the switch exceeds the maximum number of spanning-tree instances, we recommend that you configure the IEEE 802.1s Multiple STP (MSTP) on your switch to map multiple VLANs to a single spanning-tree instance. For more information about MSTP, see Chapter 21, "Configuring MSTP."

- Each routed port on the switch creates an internal VLAN for its use. These internal VLANs use extended-range VLAN numbers, and the internal VLAN ID cannot be used for an extended-range VLAN. If you try to create an extended-range VLAN with a VLAN ID that is already allocated as an internal VLAN, an error message is generated, and the command is rejected.
  - Because internal VLAN IDs are in the lower part of the extended range, we recommend that you create extended-range VLANs beginning from the highest number (4096) and moving to the lowest (1006) to reduce the possibility of using an internal VLAN ID.
  - Before configuring extended-range VLANs, enter the **show vlan internal usage** privileged EXEC command to see which VLANs have been allocated as internal VLANs.
  - If necessary, you can shut down the routed port assigned to the internal VLAN, which frees up the internal VLAN, and then create the extended-range VLAN and re-enable the port, which then uses another VLAN as its internal VLAN. See the "Creating an Extended-Range VLAN with an Internal VLAN ID" section on page 17-18.
- Although the switch supports a total of 1005 (normal-range and extended-range) VLANs, the number of routed ports, SVIs, and other configured features affects the use of the switch hardware. If you try to create an extended-range VLAN and there are not enough hardware resources available, an error message is generated, and the extended-range VLAN is rejected.

VLAN Trunks

Trunking Overview

A trunk is a point-to-point link between one or more Ethernet switch interfaces and another networking device such as a router or a switch. Ethernet trunks carry the traffic of multiple VLANs over a single link, and you can extend the VLANs across an entire network.

You can configure a trunk on a single Ethernet interface or on an EtherChannel bundle. For more information about EtherChannel, see Chapter 40, "Configuring EtherChannels."

Ethernet trunk interfaces support different trunking modes (see Table 17-3). You can set an interface as trunking or nontrunking or to negotiate trunking with the neighboring interface. To autonegotiate trunking, the interfaces must be in the same VTP domain.

Trunk negotiation is managed by the Dynamic Trunking Protocol (DTP), which is a Point-to-Point Protocol. However, some internetworking devices might forward DTP frames improperly, which could cause misconfigurations.

To avoid this, you should configure interfaces connected to devices that do not support DTP to not forward DTP frames, that is, to turn off DTP.

- If you do not intend to trunk across those links, use the **switchport mode access** interface configuration command to disable trunking.
- To enable trunking to a device that does not support DTP, use the **switchport mode trunk** and **switchport nonegotiate** interface configuration commands to cause the interface to become a trunk but to not generate DTP frames.

Table 17-3 Layer 2 Interface Modes

Mode	Function
switchport mode access	Puts the interface (access port) into permanent nontrunking mode and negotiates to convert the link into a nontrunk link. The interface becomes a nontrunk interface regardless of whether or not the neighboring interface is a trunk interface.
switchport mode dynamic auto	Makes the interface able to convert the link to a trunk link. The interface becomes a trunk interface if the neighboring interface is set to *trunk* or *desirable* mode. The default switch port mode for all Ethernet interfaces is dynamic auto.
switchport mode dynamic desirable	Makes the interface actively attempt to convert the link to a trunk link. The interface becomes a trunk interface if the neighboring interface is set to *trunk*, *desirable*, or *auto* mode.
switchport mode trunk	Puts the interface into permanent trunking mode and negotiates to convert the neighboring link into a trunk link. The interface becomes a trunk interface even if the neighboring interface is not a trunk interface.
switchport nonegotiate	Prevents the interface from generating DTP frames. You can use this command only when the interface switchport mode is access or trunk. You must manually configure the neighboring interface as a trunk interface to establish a trunk link.

IEEE 802.1Q Configuration Guidelines

The IEEE 802.1Q trunks impose these restrictions on the trunking strategy for a network:

- In a network of Cisco switches connected through IEEE 802.1Q trunks, the switches maintain one spanning-tree instance for each VLAN allowed on the trunks. Non-Cisco devices might support one spanning-tree instance for all VLANs.

  When you connect a Cisco switch to a non-Cisco device through an IEEE 802.1Q trunk, the Cisco switch combines the spanning-tree instance of the VLAN of the trunk with the spanning-tree instance of the non-Cisco IEEE 802.1Q switch. However, spanning-tree information for each VLAN is maintained by Cisco switches separated by a cloud of non-Cisco IEEE 802.1Q switches. The non-Cisco IEEE 802.1Q cloud separating the Cisco switches is treated as a single trunk link between the switches.
- Make sure the native VLAN for an IEEE 802.1Q trunk is the same on both ends of the trunk link. If the native VLAN on one end of the trunk is different from the native VLAN on the other end, spanning-tree loops might result.
- Disabling spanning tree on the native VLAN of an IEEE 802.1Q trunk without disabling spanning tree on every VLAN in the network can potentially cause spanning-tree loops. We recommend that you leave spanning tree enabled on the native VLAN of an IEEE 802.1Q trunk or disable spanning tree on every VLAN in the network. Make sure your network is loop-free before you disable spanning tree.

Default Layer 2 Ethernet Interface VLAN Settings

Table 17-4 Default Layer 2 Ethernet Interface VLAN Settings

Feature	Default Setting
Interface mode	switchport mode dynamic auto
Allowed VLAN range	VLANs 1 to 4096
VLAN range eligible for pruning	VLANs 2 to 1001
Default VLAN (for access ports)	VLAN 1
Native VLAN (for IEEE 802.1Q trunks)	VLAN 1

Ethernet Interface as a Trunk Port

Because trunk ports send and receive VTP advertisements, to use VTP you must ensure that at least one trunk port is configured on the switch and that this trunk port is connected to the trunk port of a second switch. Otherwise, the switch cannot receive any VTP advertisements.

Note By default, an interface is in Layer 2 mode. The default mode for Layer 2 interfaces is **switchport mode dynamic auto**. If the neighboring interface supports trunking and is configured to allow trunking, the link is a Layer 2 trunk or, if the interface is in Layer 3 mode, it becomes a Layer 2 trunk when you enter the **switchport** interface configuration command.

Trunking Interaction with Other Features

Trunking interacts with other features in these ways:

- A trunk port cannot be a secure port.
- A trunk port cannot be a tunnel port.
- Trunk ports can be grouped into EtherChannel port groups, but all trunks in the group must have the same configuration. When a group is first created, all ports follow the parameters set for the first port to be added to the group. If you change the configuration of one of these parameters, the switch propagates the setting you entered to all ports in the group:
  - Allowed-VLAN list.
  - STP port priority for each VLAN.
  - STP Port Fast setting.
  - Trunk status. If one port in a port group ceases to be a trunk, all ports cease to be trunks.
- We recommend that you configure no more than 24 trunk ports in PVST mode and no more than 40 trunk ports in MST mode.
- If you try to enable IEEE 802.1x on a trunk port, an error message appears, and IEEE 802.1x is not enabled. If you try to change the mode of an IEEE 802.1x-enabled port to trunk, the port mode is not changed.
- A port in dynamic mode can negotiate with its neighbor to become a trunk port. If you try to enable IEEE 802.1x on a dynamic port, an error message appears, and IEEE 802.1x is not enabled. If you try to change the mode of an IEEE 802.1x-enabled port to dynamic, the port mode is not changed.

Allowed VLANs on a Trunk

By default, a trunk port sends traffic to and receives traffic from all VLANs. All VLAN IDs, 1 to 4096, are allowed on each trunk. However, you can remove VLANs from the allowed list, preventing traffic from those VLANs from passing over the trunk. To restrict the traffic a trunk carries, use the **switchport trunk allowed vlan remove** *vlan-list* interface configuration command to remove specific VLANs from the allowed list.

Note VLAN 1 is the default VLAN on all trunk ports in all Cisco switches, and it has previously been a requirement that VLAN 1 always be enabled on every trunk link. You can use the VLAN 1 minimization feature to disable VLAN 1 on any individual VLAN trunk link so that no user traffic (including spanning-tree advertisements) is sent or received on VLAN 1.

To reduce the risk of spanning-tree loops or storms, you can disable VLAN 1 on any individual VLAN trunk port by removing VLAN 1 from the allowed list. When you remove VLAN 1 from a trunk port, the interface continues to send and receive management traffic, for example, Cisco Discovery Protocol (CDP), Port Aggregation Protocol (PAgP), Link Aggregation Control Protocol (LACP), DTP, and VTP in VLAN 1.

If a trunk port with VLAN 1 disabled is converted to a nontrunk port, it is added to the access VLAN. If the access VLAN is set to 1, the port will be added to VLAN 1, regardless of the **switchport trunk allowed** setting. The same situation applies for any VLAN that has been disabled on the port.

A trunk port can become a member of a VLAN if the VLAN is enabled, if VTP knows of the VLAN, and if the VLAN is in the allowed list for the port. When VTP detects a newly enabled VLAN and the VLAN is in the allowed list for a trunk port, the trunk port automatically becomes a member of the enabled VLAN. When VTP detects a new VLAN and the VLAN is not in the allowed list for a trunk port, the trunk port does not become a member of the new VLAN.

Native VLAN for Untagged Traffic

A trunk port configured with IEEE 802.1Q tagging can receive both tagged and untagged traffic. By default, the switch forwards untagged traffic in the native VLAN configured for the port. The native VLAN is VLAN 1 by default.

Note The native VLAN can be assigned any VLAN ID.

For information about IEEE 802.1Q configuration issues, see the "IEEE 802.1Q Configuration Guidelines" section on page 17-10.

Load Sharing Using Trunk Ports

Load sharing divides the bandwidth supplied by parallel trunks connecting switches. To avoid loops, STP normally blocks all but one parallel link between switches. Using load sharing, you divide the traffic between the links according to which VLAN the traffic belongs.

You configure load sharing on trunk ports by using STP port priorities or STP path costs. For load sharing using STP port priorities, both load-sharing links must be connected to the same switch. For load sharing using STP path costs, each load-sharing link can be connected to the same switch or to two different switches. For more information about STP, see Chapter 20, "Configuring STP."

Load Sharing Using STP Port Priorities

When two ports on the same switch form a loop, the switch uses the STP port priority to decide which port is enabled and which port is in a blocking state. You can set the priorities on a parallel trunk port so that the port carries all the traffic for a given VLAN. The trunk port with the higher priority (lower values) for a VLAN is forwarding traffic for that VLAN. The trunk port with the lower priority (higher values) for the same VLAN remains in a blocking state for that VLAN. One trunk port sends or receives all traffic for the VLAN.

Figure 17-2 shows two trunks connecting supported switches. In this example, the switches are configured as follows:

- VLANs 8 through 10 are assigned a port priority of 16 on Trunk 1.
- VLANs 3 through 6 retain the default port priority of 128 on Trunk 1.
- VLANs 3 through 6 are assigned a port priority of 16 on Trunk 2.
- VLANs 8 through 10 retain the default port priority of 128 on Trunk 2.

In this way, Trunk 1 carries traffic for VLANs 8 through 10, and Trunk 2 carries traffic for VLANs 3 through 6. If the active trunk fails, the trunk with the lower priority takes over and carries the traffic for all of the VLANs. No duplication of traffic occurs over any trunk port.

Figure 17-2 Load Sharing by Using STP Port Priorities

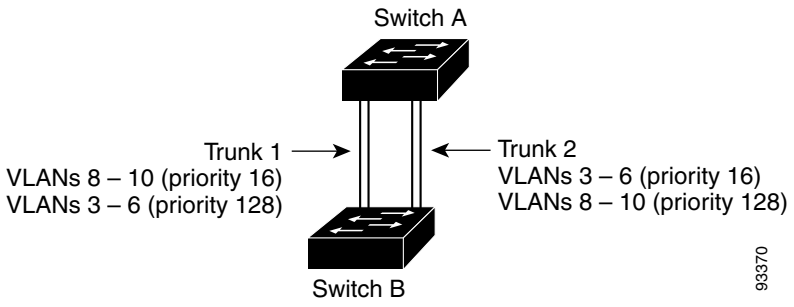

Load Sharing Using STP Path Cost

You can configure parallel trunks to share VLAN traffic by setting different path costs on a trunk and associating the path costs with different sets of VLANs, blocking different ports for different VLANs. The VLANs keep the traffic separate and maintain redundancy in the event of a lost link.

In Figure 17-3, Trunk ports 1 and 2 are configured as 100BASE-T ports. These VLAN path costs are assigned:

- VLANs 2 through 4 are assigned a path cost of 30 on Trunk port 1.
- VLANs 8 through 10 retain the default 100BASE-T path cost on Trunk port 1 of 19.
- VLANs 8 through 10 are assigned a path cost of 30 on Trunk port 2.
- VLANs 2 through 4 retain the default 100BASE-T path cost on Trunk port 2 of 19.

Figure 17-3 Load-Sharing Trunks with Traffic Distributed by Path Cost

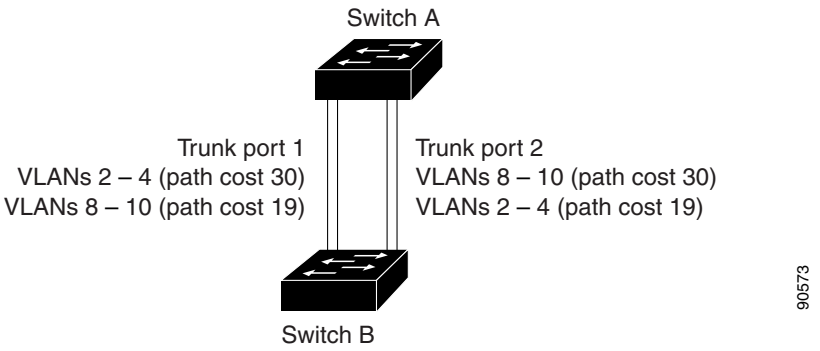

See the "Configuring Load Sharing Using STP Path Cost" section on page 17-21.

VMPS

The VLAN Query Protocol (VQP) is used to support dynamic-access ports, which are not permanently assigned to a VLAN, but give VLAN assignments based on the MAC source addresses seen on the port. Each time an unknown MAC address is seen, the switch sends a VQP query to a remote VMPS; the query includes the newly seen MAC address and the port on which it was seen. The VMPS responds with a VLAN assignment for the port. The switch cannot be a VMPS server but can act as a client to the VMPS and communicate with it through VQP.

Each time the client switch receives the MAC address of a new host, it sends a VQP query to the VMPS. When the VMPS receives this query, it searches its database for a MAC-address-to-VLAN mapping. The server response is based on this mapping and whether or not the server is in open or secure mode. In secure mode, the server shuts down the port when an illegal host is detected. In open mode, the server simply denies the host access to the port.

If the port is currently *unassigned* (that is, it does not yet have a VLAN assignment), the VMPS provides one of these responses:

- If the host is allowed on the port, the VMPS sends the client a *vlan-assignment* response containing the assigned VLAN name and allowing access to the host.
- If the host is not allowed on the port and the VMPS is in open mode, the VMPS sends an *access-denied* response.
- If the VLAN is not allowed on the port and the VMPS is in secure mode, the VMPS sends a *port-shutdown* response.

If the port already has a VLAN assignment, the VMPS provides one of these responses:

- If the VLAN in the database matches the current VLAN on the port, the VMPS sends a *success* response, allowing access to the host.
- If the VLAN in the database does not match the current VLAN on the port and active hosts exist on the port, the VMPS sends an *access-denied* or a *port-shutdown* response, depending on the secure mode of the VMPS.

If the switch receives an *access-denied* response from the VMPS, it continues to block traffic to and from the host MAC address. The switch continues to monitor the packets directed to the port and sends a query to the VMPS when it identifies a new host address. If the switch receives a *port-shutdown* response from the VMPS, it disables the port. The port must be manually reenabled by using Network Assistant, the CLI, or SNMP.

Dynamic-Access Port VLAN Membership

A dynamic-access port can belong to only one VLAN with an ID from 1 to 4096. When the link comes up, the switch does not forward traffic to or from this port until the VMPS provides the VLAN assignment. The VMPS receives the source MAC address from the first packet of a new host connected to the dynamic-access port and attempts to match the MAC address to a VLAN in the VMPS database.

If there is a match, the VMPS sends the VLAN number for that port. If the client switch was not previously configured, it uses the domain name from the first VTP packet it receives on its trunk port from the VMPS. If the client switch was previously configured, it includes its domain name in the query packet to the VMPS to obtain its VLAN number. The VMPS verifies that the domain name in the packet matches its own domain name before accepting the request and responds to the client with the assigned VLAN number for the client. If there is no match, the VMPS either denies the request or shuts down the port (depending on the VMPS secure mode setting).

Multiple hosts (MAC addresses) can be active on a dynamic-access port if they are all in the same VLAN; however, the VMPS shuts down a dynamic-access port if more than 20 hosts are active on the port.

If the link goes down on a dynamic-access port, the port returns to an isolated state and does not belong to a VLAN. Any hosts that come online through the port are checked again through the VQP with the VMPS before the port is assigned to a VLAN.

Dynamic-access ports can be used for direct host connections, or they can connect to a network. A maximum of 20 MAC addresses are allowed per port on the switch. A dynamic-access port can belong to only one VLAN at a time, but the VLAN can change over time, depending on the MAC addresses seen.

Default VMPS Client Settings

Table 17-5 Default VMPS Client and Dynamic-Access Port Settings

Feature	Default Setting
VMPS domain server	None
VMPS reconfirm interval	60 minutes
VMPS server retry count	3
Dynamic-access ports	None configured

VMPS Configuration Guidelines

These guidelines and restrictions apply to dynamic-access port VLAN membership:

- You should configure the VMPS before you configure ports as dynamic-access ports.
- When you configure a port as a dynamic-access port, the spanning-tree Port Fast feature is automatically enabled for that port. The Port Fast mode accelerates the process of bringing the port into the forwarding state.
- IEEE 802.1x ports cannot be configured as dynamic-access ports. If you try to enable IEEE 802.1x on a dynamic-access (VQP) port, an error message appears, and IEEE 802.1x is not enabled. If you try to change an IEEE 802.1x-enabled port to dynamic VLAN assignment, an error message appears, and the VLAN configuration is not changed.

- Trunk ports cannot be dynamic-access ports, but you can enter the **switchport access vlan dynamic** interface configuration command for a trunk port. In this case, the switch retains the setting and applies it if the port is later configured as an access port.

  You must turn off trunking on the port before the dynamic-access setting takes effect.
- Dynamic-access ports cannot be monitor ports.
- Secure ports cannot be dynamic-access ports. You must disable port security on a port before it becomes dynamic.
- Private VLAN ports cannot be dynamic-access ports.
- Dynamic-access ports cannot be members of an EtherChannel group.
- Port channels cannot be configured as dynamic-access ports.
- A dynamic-access port can participate in fallback bridging.
- The VTP management domain of the VMPS client and the VMPS server must be the same.
- The VLAN configured on the VMPS server should not be a voice VLAN.

VMPS Reconfirmation Interval

VMPS clients periodically reconfirm the VLAN membership information received from the VMPS.You can set the number of minutes after which reconfirmation occurs.

If you are configuring a member switch in a cluster, this parameter must be equal to or greater than the reconfirmation setting on the command switch. You must also first use the **rcommand** privileged EXEC command to log in to the member switch.

Dynamic-Access Port VLAN Membership

The VMPS shuts down a dynamic-access port under these conditions:

- The VMPS is in secure mode, and it does not allow the host to connect to the port. The VMPS shuts down the port to prevent the host from connecting to the network.
- More than 20 active hosts reside on a dynamic-access port.

To reenable a disabled dynamic-access port, enter the **shutdown** interface configuration command followed by the **no shutdown** interface configuration command.

How to Configure VLANs

Creating or Modifying an Ethernet VLAN

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**vlan** *vlan-id*	Enters a VLAN ID, and enters VLAN configuration mode. **Note** The available VLAN ID range for this command is 1 to 4096. For information about adding VLAN IDs greater than 1005 (extended-range VLANs), see the "Creating an Extended-Range VLAN" section on page 17-18.
Step 3	**name** *vlan-name*	(Optional) Enters a name for the VLAN. If no name is entered for the VLAN, the default is to append the *vlan-id* with leading zeros to the word VLAN. For example, VLAN0004 is a default VLAN name for VLAN 4.
Step 4	**mtu** *mtu-size*	(Optional) Changes the MTU size (or other VLAN characteristic).
Step 5	**remote-span**	(Optional) Configures the VLAN as the RSPAN VLAN for a remote SPAN session. **Note** For more information on remote SPAN, see Chapter 30, "Configuring SPAN and RSPAN."
Step 6	**end**	Returns to privileged EXEC mode.

Deleting a VLAN

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**no vlan** *vlan-id*	Removes the VLAN by entering the VLAN ID.
Step 3	**end**	Returns to privileged EXEC mode.

Assigning Static-Access Ports to a VLAN

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode
Step 2	**interface** *interface-id*	Enters the interface to be added to the VLAN.
Step 3	**switchport mode access**	Defines the VLAN membership mode for the port (Layer 2 access port).
Step 4	**switchport access vlan** *vlan-id*	Assigns the port to a VLAN. Valid VLAN IDs are 1 to 4096.
Step 5	**end**	Returns to privileged EXEC mode.

Creating an Extended-Range VLAN

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**vtp mode transparent**	Configures the switch for VTP transparent mode and disables VTP. **Note** This step is not required for VTP version 3.
Step 3	**vlan** *vlan-id*	Enters an extended-range VLAN ID and enters VLAN configuration mode. The range is 1006 to 4096.
Step 4	**mtu** *mtu-size*	(Optional) Modifies the VLAN by changing the MTU size. **Note** Although all VLAN commands appear in the CLI help, only the **mtu** *mtu-size*, **private-vlan**, and **remote-span** commands are supported for extended-range VLANs.
Step 5	**remote-span**	(Optional) Configures the VLAN as the RSPAN VLAN. See the "Configuring a VLAN as an RSPAN VLAN" section on page 30-14.
Step 6	**end**	Returns to privileged EXEC mode.

Creating an Extended-Range VLAN with an Internal VLAN ID

	Command	Purpose
Step 1	**show vlan internal usage**	Displays the VLAN IDs being used internally by the switch. If the VLAN ID that you want to use is an internal VLAN, the display shows the routed port that is using the VLAN ID. Enter that port number in Step 3.
Step 2	**configure terminal**	Enters global configuration mode.
Step 3	**interface** *interface-id*	Specifies the interface ID for the routed port that is using the VLAN ID, and enters interface configuration mode.
Step 4	**shutdown**	Shuts down the port to free the internal VLAN ID.
Step 5	**exit**	Returns to global configuration mode.
Step 6	**vtp mode transparent**	Sets the VTP mode to transparent for creating extended-range VLANs. **Note** This step is not required for VTP version 3.
Step 7	**vlan** *vlan-id*	Enters the new extended-range VLAN ID, and enters VLAN configuration mode.
Step 8	**exit**	Exits from VLAN configuration mode, and returns to global configuration mode.
Step 9	**interface** *interface-id*	Specifies the interface ID for the routed port that you shut down in Step 4, and enters interface configuration mode.
Step 10	**no shutdown**	Reenables the routed port. It will be assigned a new internal VLAN ID.
Step 11	**end**	Returns to privileged EXEC mode.

Configuring an Ethernet Interface as a Trunk Port

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the port to be configured for trunking, and enters interface configuration mode.
Step 3	**switchport mode** {**dynamic** {**auto** | **desirable**} | **trunk**}	Configures the interface as a Layer 2 trunk (required only if the interface is a Layer 2 access port or tunnel port or to specify the trunking mode). • **dynamic auto**—Sets the interface to a trunk link if the neighboring interface is set to trunk or desirable mode. This is the default. • **dynamic desirable**—Sets the interface to a trunk link if the neighboring interface is set to trunk, desirable, or auto mode. • **trunk**—Sets the interface in permanent trunking mode and negotiate to convert the link to a trunk link even if the neighboring interface is not a trunk interface.
Step 4	**switchport access vlan** *vlan-id*	(Optional) Specifies the default VLAN, which is used if the interface stops trunking.
Step 5	**switchport trunk native vlan** *vlan-id*	Specifies the native VLAN for IEEE 802.1Q trunks.
Step 6	**end**	Returns to privileged EXEC mode.

Defining the Allowed VLANs on a Trunk

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the port to be configured, and enters interface configuration mode.
Step 3	**switchport mode trunk**	Configures the interface as a VLAN trunk port.
Step 4	**switchport trunk allowed vlan** {**add** | **all** | **except** | **remove**} *vlan-list*	(Optional) Configures the list of VLANs allowed on the trunk.
Step 5	**end**	Returns to privileged EXEC mode.

Changing the Pruning-Eligible List

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Selects the trunk port for which VLANs should be pruned, and enters interface configuration mode.

	Command	Purpose
Step 3	**switchport trunk pruning vlan** {**add** | **except** | **none** | **remove**} *vlan-list* [,*vlan*[,*vlan*[,,]]	Configures the list of VLANs allowed to be pruned from the trunk. (See the "VTP Pruning" section on page 18-7.)
Step 4	**end**	Returns to privileged EXEC mode.

Configuring the Native VLAN for Untagged Traffic

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Defines the interface that is configured as the IEEE 802.1Q trunk, and enters interface configuration mode.
Step 3	**switchport trunk native vlan** *vlan-id*	Configures the VLAN that is sending and receiving untagged traffic on the trunk port.
Step 4	**end**	Returns to privileged EXEC mode.

Load Sharing Using STP Port Priorities

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode on Switch A.
Step 2	**vtp domain** *domain-name*	Configures a VTP administrative domain. The domain name can be 1 to 32 characters.
Step 3	**vtp mode server**	Configures Switch A as the VTP server.
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show vtp status**	Verifies the VTP configuration on both Switch A and Switch B.
Step 6	**show vlan**	Verifies that the VLANs exist in the database on Switch A.
Step 7	**configure terminal**	Enters global configuration mode.
Step 8	**interface** *interface-id_1*	Defines the interface to be configured as a trunk, and enters interface configuration mode.
Step 9	**switchport mode trunk**	Configures the port as a trunk port.
Step 10	**end**	Returns to privileged EXEC mode.
Step 11	**show interfaces** *interface-id_1* **switchport**	Verifes the VLAN configuration.
Step 12	Repeat Steps 7 through 10 on Switch A for a second port in the switch.	
Step 13	Repeat Steps 7 through 10 on Switch B to configure the trunk ports that connect to the trunk ports configured on Switch A.	
Step 14	**show vlan**	When the trunk links come up, VTP passes the VTP and VLAN information to Switch B. Verifies that Switch B has learned the VLAN configuration.
Step 15	**configure terminal**	Enters global configuration mode on Switch A.
Step 16	**interface** *interface-id_1*	Defines the interface to set the STP port priority, and enters interface configuration mode.
Step 17	**spanning-tree vlan 8-10 port-priority 16**	Assigns the port priority of 16 for VLANs 8 through 10.
Step 18	**exit**	Returns to global configuration mode.
Step 19	**interface** *interface-id_2*	Defines the interface to set the STP port priority, and enters interface configuration mode.
Step 20	**spanning-tree vlan 3-6 port-priority 16**	Assigns the port priority of 16 for VLANs 3 through 6.
Step 21	**end**	Returns to privileged EXEC mode.

Configuring Load Sharing Using STP Path Cost

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode on Switch A.
Step 2	**interface** *interface-id_1*	Defines the interface to be configured as a trunk, and enters interface configuration mode.
Step 3	**switchport mode trunk**	Configures the port as a trunk port.

	Command	Purpose
Step 4	**exit**	Returns to global configuration mode.
Step 5		Repeat Steps 2 through 4 on a second interface in Switch A.
Step 6	**end**	Returns to privileged EXEC mode.
Step 7	**show running-config**	Verifies your entries. In the display, make sure that the interfaces are configured as trunk ports.
Step 8	**show vlan**	When the trunk links come up, Switch A receives the VTP information from the other switches. Verifies that Switch A has learned the VLAN configuration.
Step 9	**configure terminal**	Enters global configuration mode.
Step 10	**interface** *interface-id_1*	Defines the interface on which to set the STP cost, and enters interface configuration mode.
Step 11	**spanning-tree vlan 2-4 cost 30**	Sets the spanning-tree path cost to 30 for VLANs 2 through 4.
Step 12	**end**	Returns to global configuration mode.
Step 13	Repeat Steps 9 through 12 on the other configured trunk interface on Switch A, and set the spanning-tree path cost to 30 for VLANs 8, 9, and 10.	
Step 14	**exit**	Returns to privileged EXEC mode.
Step 15	**show running-config**	Verifies your entries. In the display, verify that the path costs are set correctly for both trunk interfaces.

Configuring the VMPS Client

You configure dynamic VLANs by using the VMPS (VLAN Membership Policy Server). The switch can be a VMPS client; it cannot be a VMPS server.

Entering the IP Address of the VMPS

Before You Begin

- You must first enter the IP address of the server to configure the switch as a client.
- You must have IP connectivity to the VMPS for dynamic-access ports to work. You can test for IP connectivity by pinging the IP address of the VMPS and verifying that you get a response.
- If the VMPS is being defined for a cluster of switches, enter the address on the command switch.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**vmps server** *ipaddress* **primary**	Enters the IP address of the switch acting as the primary VMPS server.
Step 3	**vmps server** *ipaddress*	(Optional) Enters the IP address of the switch acting as a secondary VMPS server. You can enter up to three secondary server addresses.
Step 4	**vmps reconfirm**	(Optional) Reconfirms dynamic-access port VLAN membership.

	Command	Purpose
Step 5	**vmps retry** *count*	(Optional) Changes the retry count.
Step 6	**end**	Returns to privileged EXEC mode.

Configuring Dynamic-Access Ports on VMPS Clients

Before You Begin

If you are configuring a port on a cluster member switch as a dynamic-access port, first use the **rcommand** privileged EXEC command to log in to the cluster member switch.

Caution Dynamic-access port VLAN membership is for end stations or hubs connected to end stations. Connecting dynamic-access ports to other switches can cause a loss of connectivity.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the switch port that is connected to the end station, and enters interface configuration mode.
Step 3	**switchport mode access**	Sets the port to access mode.
Step 4	**switchport access vlan dynamic**	Configures the port as eligible for dynamic VLAN membership. The dynamic-access port must be connected to an end station.
Step 5	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining VLANs

Command	Purpose
copy running-config startup config	Saves your entries in the configuration file • To save an extended-range VLAN configuration, you need to save the VTP transparent mode configuration and the extended-range VLAN configuration in the switch startup configuration file. Otherwise, if the switch resets, it will default to VTP server mode, and the extended-range VLAN IDs will not be saved. • This step is not required for VTP version 3 because VLANs are saved in the VLAN database.
show interfaces *interface-id* **switchport**	Displays the switch port configuration of the interface.
show interfaces *interface-id* **trunk**	Displays the trunk configuration of the interface.
show running-config interface *interface-id*	Verifies the VLAN membership mode of the interface.
show vmps	Verifies your VMPS entries.
show vlan	Verifies your VLAN entries.

Configuration Examples for Configuring VLANs

VMPS Network: Example

Figure 17-4 shows a network with a VMPS server switch and VMPS client switches with dynamic-access ports. In this example, these assumptions apply:

- The VMPS server and the VMPS client are separate switches.
- The Catalyst 6500 series Switch A is the primary VMPS server.
- The Catalyst 6500 series Switch C and Switch J are secondary VMPS servers.
- End stations are connected to the clients, Switch B and Switch I.
- The database configuration file is stored on the TFTP server with the IP address 172.20.22.7.

Figure 17-4 Dynamic Port VLAN Membership Configuration

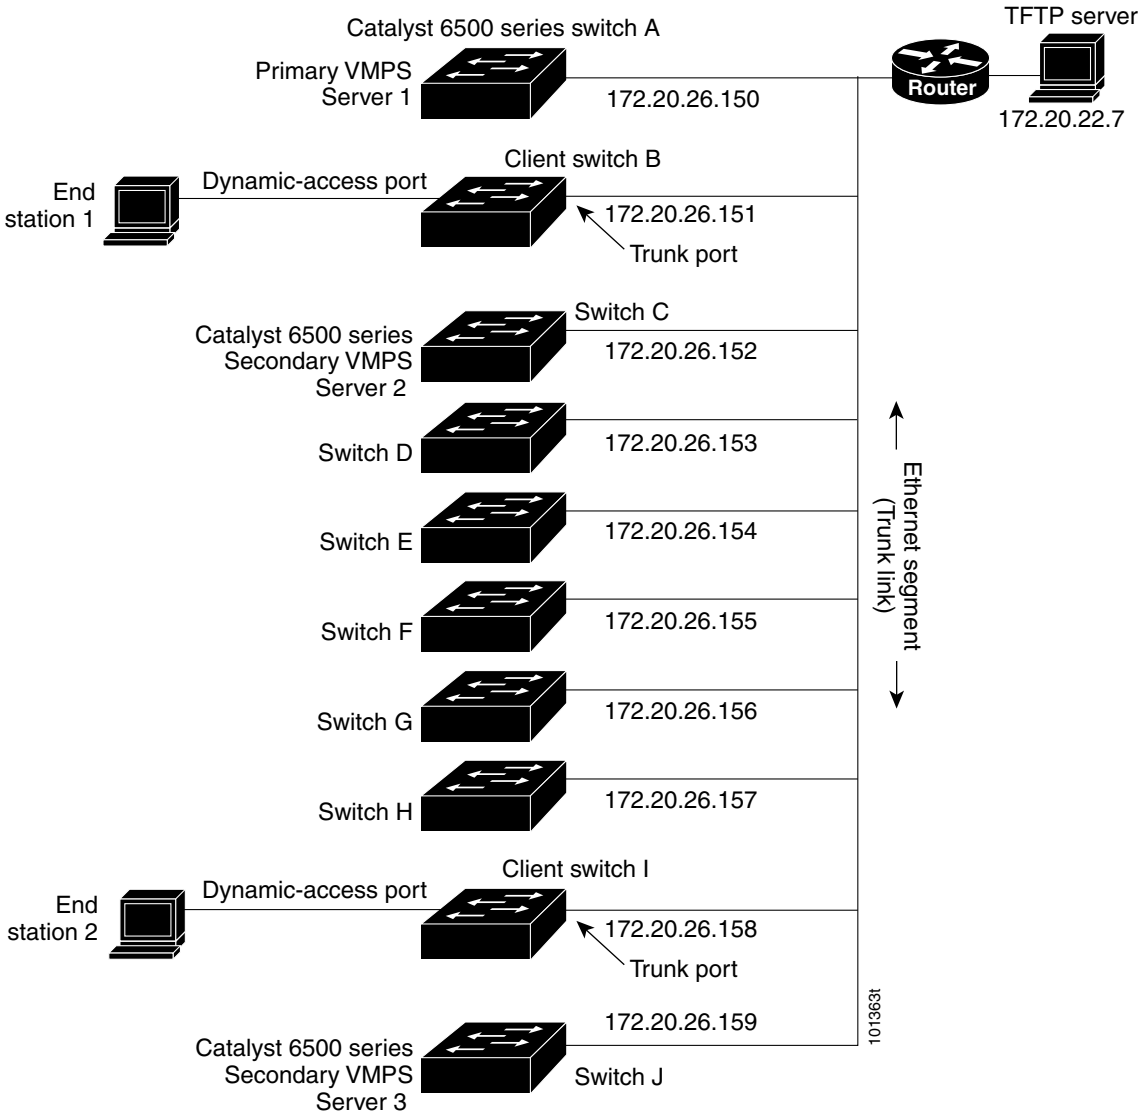

Configuring a VLAN: Example

This example shows how to create Ethernet VLAN 20, name it *test20,* and add it to the VLAN database:

```
Switch# configure terminal
Switch(config)# vlan 20
Switch(config-vlan)# name test20
Switch(config-vlan)# end
```

Configuring an Access Port in a VLAN: Example

This example shows how to configure a port as an access port in VLAN 2:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan 2
Switch(config-if)# end
```

Configuring an Extended-Range VLAN: Example

This example shows how to create a new extended-range VLAN with all default characteristics:

```
Switch(config)# vtp mode transparent
Switch(config)# vlan 2000
Switch(config-vlan)# end
Switch# copy running-config startup config
```

Configuring a Trunk Port: Example

This example shows how to configure a port as an IEEE 802.1Q trunk. The example assumes that the neighbor interface is configured to support IEEE 802.1Q trunking.

```
Switch(config)# interface gigabitethernet1/2
Switch(config-if)# switchport mode dynamic desirable
Switch(config-if)# end
```

Removing a VLAN: Example

This example shows how to remove VLAN 2 from the allowed VLAN list on a port:

```
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# switchport trunk allowed vlan remove 2
Switch(config-if)# end
```

Show VMPS Output: Example

This is an example of output for the **show vmps** privileged EXEC command:

```
Switch# show vmps
VQP Client Status:
--------------------
VMPS VQP Version:   1
```

```
Reconfirm Interval: 60 min
Server Retry Count: 3
VMPS domain server: 172.20.128.86 (primary, current)
                    172.20.128.87

Reconfirmation status
---------------------
VMPS Action:         other
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
VTP pruning configuration	*Chapter 18, "Configuring VTP"*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

CHAPTER 18

Configuring VTP

Finding VTP Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for Configuring VTP

- When you configure VTP, you must configure a trunk port so that the switch can send and receive VTP advertisements to and from other switches in the domain. For more information, see the "Configuring an Ethernet Interface as a Trunk Port" section on page 17-19.
- Before adding a VTP client switch to a VTP domain, always verify that its VTP configuration revision number is *lower* than the configuration revision number of the other switches in the VTP domain. Switches in a VTP domain always use the VLAN configuration of the switch with the highest VTP configuration revision number. If you add a switch that has a revision number higher than the revision number in the VTP domain, it can erase all VLAN information from the VTP server and VTP domain. See the "Adding a VTP Client Switch to a VTP Domain" section on page 18-13 for the procedure for verifying and resetting the VTP configuration revision number.

Restrictions for Configuring VTP

- For VTP version 3, the switch must be running the LAN Base image.
- VTP version 1 and VTP version 2 are not interoperable on switches in the same VTP domain. Do not enable VTP version 2 unless every switch in the VTP domain supports version 2.
- In VTP versions 1 and 2, when you configure extended-range VLANs on the switch, the switch must be in VTP transparent mode. VTP version 3 also supports creating extended-range VLANs in client or server mode.

Information About Configuring VTP

VTP

A VLAN Trunking Protocol (VTP) is a Layer 2 messaging protocol that maintains VLAN configuration consistency by managing the addition, deletion, and renaming of VLANs on a network-wide basis. VTP minimizes misconfigurations and configuration inconsistencies that can cause several problems, such as duplicate VLAN names, incorrect VLAN-type specifications, and security violations.

Before you create VLANs, you must decide whether to use VTP in your network. Using VTP, you can make configuration changes centrally on one or more switches and have those changes automatically communicated to all the other switches in the network. Without VTP, you cannot send information about VLANs to other switches.

VTP is designed to work in an environment where updates are made on a single switch and are sent through VTP to other switches in the domain. It does not work well in a situation where multiple updates to the VLAN database occur simultaneously on switches in the same domain, which would result in an inconsistency in the VLAN database.

The switch supports 1005 VLANs, but the number of configured features affects the usage of the switch hardware. If the switch is notified by VTP of a new VLAN and the switch is already using the maximum available hardware resources, it sends a message that there are not enough hardware resources available and shuts down the VLAN. The output of the **show vlan** user EXEC command shows the VLAN in a suspended state.

VTP version 1 and version 2 support only normal-range VLANs (VLAN IDs 1 to 1005). VTP version 3 supports the entire VLAN range (VLANs 1 to 4096). Extended range VLANs (VLANs 1006 to 4096) are supported only in VTP version 3. You cannot convert from VTP version 3 to VTP version 2 if extended VLANs are configured in the domain.

VTP Domain

A VTP domain (also called a VLAN management domain) consists of one switch or several interconnected switches under the same administrative responsibility sharing the same VTP domain name. A switch can be in only one VTP domain. You make global VLAN configuration changes for the domain.

By default, the switch is in the VTP no-management-domain state until it receives an advertisement for a domain over a trunk link (a link that carries the traffic of multiple VLANs) or until you configure a domain name. Until the management domain name is specified or learned, you cannot create or modify VLANs on a VTP server, and VLAN information is not propagated over the network.

If the switch receives a VTP advertisement over a trunk link, it inherits the management domain name and the VTP configuration revision number. The switch then ignores advertisements with a different domain name or an earlier configuration revision number.

When you make a change to the VLAN configuration on a VTP server, the change is propagated to all switches in the VTP domain. VTP advertisements are sent over all IEEE trunk connections, including IEEE 802.1Q. VTP dynamically maps VLANs with unique names and internal index associates across multiple LAN types. Mapping eliminates excessive device administration required from network administrators.

If you configure a switch for VTP transparent mode, you can create and modify VLANs, but the changes are not sent to other switches in the domain, and they affect only the individual switch. However, configuration changes made when the switch is in this mode are saved in the switch running configuration and can be saved to the switch startup configuration file.

For domain name and password configuration guidelines, see the "VTP Configuration Guidelines" section on page 18-9.

VTP Modes

Table 18-1 VTP Modes

VTP Mode	Description
VTP server	In VTP server mode, you can create, modify, and delete VLANs, and specify other configuration parameters (such as the VTP version) for the entire VTP domain. VTP servers advertise their VLAN configurations to other switches in the same VTP domain and synchronize their VLAN configurations with other switches based on advertisements received over trunk links. VTP server is the default mode. **Note** In VTP server mode, VLAN configurations are saved in NVRAM. If the switch detects a failure while writing a configuration to NVRAM, VTP mode automatically changes from server mode to client mode. If this happens, the switch cannot be returned to VTP server mode until the NVRAM is functioning.
VTP client	A VTP client behaves like a VTP server and transmits and receives VTP updates on its trunks, but you cannot create, change, or delete VLANs on a VTP client. VLANs are configured on another switch in the domain that is in server mode. In VTP versions 1 and 2, in VTP client mode, VLAN configurations are not saved in NVRAM. In VTP version 3, VLAN configurations are saved in NVRAM in client mode.
VTP transparent	VTP transparent switches do not participate in VTP. A VTP transparent switch does not advertise its VLAN configuration and does not synchronize its VLAN configuration based on received advertisements. However, in VTP version 2 or version 3, transparent switches do forward VTP advertisements that they receive from other switches through their trunk interfaces. You can create, modify, and delete VLANs on a switch in VTP transparent mode. In VTP versions 1 and 2, the switch must be in VTP transparent mode when you create extended-range VLANs. VTP version 3 also supports creating extended-range VLANs in client or server mode. See the "Creating an Extended-Range VLAN" section on page 17-18. When the switch is in VTP transparent mode, the VTP and VLAN configurations are saved in NVRAM, but they are not advertised to other switches. In this mode, VTP mode and domain name are saved in the switch running configuration, and you can save this information in the switch startup configuration file by using the **copy running-config startup-config** privileged EXEC command.
VTP off	A switch in VTP off mode functions in the same manner as a VTP transparent switch, except that it does not forward VTP advertisements on trunks.

VTP Mode Guidelines

- For VTP version 1 and version 2, if extended-range VLANs are configured on the switch, you cannot change VTP mode to client or server. You receive an error message, and the configuration is not allowed. VTP version 1 and version 2 do not propagate configuration information for extended range VLANs (VLANs 1006 to 4096). You must manually configure these VLANs on each device.

Note For VTP version 1 and 2, before you create extended-range VLANs (VLAN IDs 1006 to 4096), you must set VTP mode to transparent by using the **vtp mode transparent** global configuration command. Save this configuration to the startup configuration so that the switch starts in VTP transparent mode. Otherwise, you lose the extended-range VLAN configuration if the switch resets and boots up in VTP server mode (the default).

- VTP version 3 supports extended-range VLANs. If extended VLANs are configured, you cannot convert from VTP version 3 to VTP version 2.
- If you configure the switch for VTP client mode, the switch does not create the VLAN database file (vlan.dat). If the switch is then powered off, it resets the VTP configuration to the default. To keep the VTP configuration with VTP client mode after the switch restarts, you must first configure the VTP domain name before the VTP mode.
- When a switch is in VTP server mode, you can change the VLAN configuration and have it propagated throughout the network.
- When a switch is in VTP client mode, you cannot change its VLAN configuration. The client switch receives VTP updates from a VTP server in the VTP domain and then modifies its configuration accordingly.
- When you configure the switch for VTP transparent mode, VTP is disabled on the switch. The switch does not send VTP updates and does not act on VTP updates received from other switches. However, a VTP transparent switch running VTP version 2 does forward received VTP advertisements on its trunk links.
- VTP off mode is the same as VTP transparent mode except that VTP advertisements are not forwarded.

Caution If all switches are operating in VTP client mode, do not configure a VTP domain name. If you do, it is impossible to make changes to the VLAN configuration of that domain. Therefore, make sure you configure at least one switch as a VTP server.

VTP Advertisements

Each switch in the VTP domain sends periodic global configuration advertisements from each trunk port to a reserved multicast address. Neighboring switches receive these advertisements and update their VTP and VLAN configurations as necessary.

VTP advertisements distribute this global domain information:

- VTP domain name
- VTP configuration revision number
- Update identity and update timestamp
- MD5 digest VLAN configuration, including maximum transmission unit (MTU) size for each VLAN
- Frame format

VTP advertisements distribute this VLAN information for each configured VLAN:

- VLAN IDs (IEEE 802.1Q)
- VLAN name

- VLAN type
- VLAN state
- Additional VLAN configuration information specific to the VLAN type

In VTP version 3, VTP advertisements also include the primary server ID, an instance number, and a start index.

VTP Version 2

If you use VTP in your network, you must decide which version of VTP to use. By default, VTP operates in version 1.

VTP version 2 supports these features that are not supported in version 1:

- Token Ring support—VTP version 2 supports Token Ring Bridge Relay Function (TrBRF) and Token Ring Concentrator Relay Function (TrCRF) VLANs. For more information about Token Ring VLANs, see the "Normal-Range VLANs" section on page 17-4.
- Unrecognized Type-Length-Value (TLV) support—A VTP server or client propagates configuration changes to its other trunks, even for TLVs it is not able to parse. The unrecognized TLV is saved in NVRAM when the switch is operating in VTP server mode.
- Version-Dependent Transparent Mode—In VTP version 1, a VTP transparent switch inspects VTP messages for the domain name and version and forwards a message only if the version and domain name match. Although VTP version 2 supports only one domain, a VTP version 2 transparent switch forwards a message only when the domain name matches.
- Consistency Checks—In VTP version 2, VLAN consistency checks (such as VLAN names and values) are performed only when you enter new information through the CLI or SNMP. Consistency checks are not performed when new information is obtained from a VTP message or when information is read from NVRAM. If the MD5 digest on a received VTP message is correct, its information is accepted.

VTP Version 3

VTP version 3 supports these features that are not supported in version 1 or version 2:

- Enhanced authentication—You can configure the authentication as **hidden** or **secret**. When **hidden**, the secret key from the password string is saved in the VLAN database file, but it does not appear in plain text in the configuration. Instead, the key associated with the password is saved in hexadecimal format in the running configuration. You must reenter the password if you enter a takeover command in the domain. When you enter the **secret** keyword, you can directly configure the password secret key.
- Support for extended range VLAN (VLANs 1006 to 4096) database propagation. VTP versions 1 and 2 propagate only VLANs 1 to 1005. If extended VLANs are configured, you cannot convert from VTP version 3 to version 1 or 2.

  Note VTP pruning still applies only to VLANs 1 to 1005, and VLANs 1002 to 1005 are still reserved and cannot be modified.

- Support for any database in a domain. In addition to propagating VTP information, version 3 can propagate Multiple Spanning Tree (MST) protocol database information. A separate instance of the VTP protocol runs for each application that uses VTP.

- VTP primary server and VTP secondary servers. A VTP primary server updates the database information and sends updates that are honored by all devices in the system. A VTP secondary server can only back up the updated VTP configurations received from the primary server to its NVRAM.

  By default, all devices come up as secondary servers. You can enter the **vtp primary** privileged EXEC command to specify a primary server. Primary server status is only needed for database updates when the administrator issues a takeover message in the domain. You can have a working VTP domain without any primary servers. Primary server status is lost if the device reloads or domain parameters change, even when a password is configured on the switch.

- The option to turn VTP on or off on a per-trunk (per-port) basis. You can enable or disable VTP per port by entering the [**no**] **vtp** interface configuration command. When you disable VTP on trunking ports, all VTP instances for that port are disabled. You cannot set VTP to *off* for the MST database and *on* for the VLAN database on the same port.

  When you globally set VTP mode to off, it applies to all the trunking ports in the system. However, you can specify on or off on a per-VTP instance basis. For example, you can configure the switch as a VTP server for the VLAN database but with VTP *off* for the MST database.

VTP Version Guidelines

Follow these guidelines when deciding which VTP version to implement:

- All switches in a VTP domain must have the same domain name, but they do not need to run the same VTP version.
- A VTP version 2-capable switch can operate in the same VTP domain as a switch running VTP version 1 if version 2 is disabled on the version 2-capable switch (version 2 is disabled by default).
- If a switch running VTP version 1 but capable of running VTP version 2 receives VTP version 3 advertisements, it automatically moves to VTP version 2.
- If a switch running VTP version 3 is connected to a switch running VTP version 1, the VTP version 1 switch moves to VTP version 2, and the VTP version 3 switch sends scaled-down versions of the VTP packets so that the VTP version 2 switch can update its database.
- A switch running VTP version 3 cannot move to version 1 or 2 if it has extended VLANs.
- Do not enable VTP version 2 on a switch unless all of the switches in the same VTP domain are version-2-capable. When you enable version 2 on a switch, all of the version-2-capable switches in the domain enable version 2. If there is a version 1-only switch, it does not exchange VTP information with switches that have version 2 enabled.
- We recommend placing VTP version 1 and 2 switches at the edge of the network because they do not forward VTP version 3 advertisements.
- If there are TrBRF and TrCRF Token Ring networks in your environment, you must enable VTP version 2 or version 3 for Token Ring VLAN switching to function properly. To run Token Ring and Token Ring-Net, disable VTP version 2.
- VTP version 1 and version 2 do not propagate configuration information for extended range VLANs (VLANs 1006 to 4096). You must configure these VLANs manually on each device. VTP version 3 supports extended-range VLANs. You cannot convert from VTP version 3 to VTP version 2 if extended VLANs are configured.
- When a VTP version 3 device trunk port receives messages from a VTP version 2 device, it sends a scaled-down version of the VLAN database on that particular trunk in VTP version 2 format. A VTP version 3 device does not send VTP version 2-formatted packets on a trunk unless it first receives VTP version 2 packets on that trunk port.

- When a VTP version 3 device detects a VTP version 2 device on a trunk port, it continues to send VTP version 3 packets, in addition to VTP version 2 packets, to allow both kinds of neighbors to coexist on the same trunk.
- A VTP version 3 device does not accept configuration information from a VTP version 2 or version 1 device.
- Two VTP version 3 regions can only communicate in transparent mode over a VTP version 1 or version 2 region.
- Devices that are only VTP version 1 capable cannot interoperate with VTP version 3 devices.
- VTP version 2 and version 3 are disabled by default.
- When you enable VTP version 2 on a switch, every VTP version 2-capable switch in the VTP domain enables version 2. To enable VTP version 3, you must manually configure it on each switch.
- With VTP versions 1 and 2, you can configure the version only on switches in VTP server or transparent mode. If a switch is running VTP version 3, you can change to version 2 when the switch is in client mode if no extended VLANs exist, no private VLANs exist, and no hidden password was configured.

Caution In VTP version 3, both the primary and secondary servers can exist on an instance in the domain.

VTP Pruning

VTP pruning increases network available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to reach the destination devices. Without VTP pruning, a switch floods broadcast, multicast, and unknown unicast traffic across all trunk links within a VTP domain even though receiving switches might discard them. VTP pruning is disabled by default.

VTP pruning blocks unneeded flooded traffic to VLANs on trunk ports that are included in the pruning-eligible list. Only VLANs included in the pruning-eligible list can be pruned. By default, VLANs 2 through 1001 are pruning eligible switch trunk ports. If the VLANs are configured as pruning-ineligible, the flooding continues. VTP pruning is supported in all VTP versions.

Figure 18-1 shows a switched network without VTP pruning enabled. Port 1 on Switch A and Port 2 on Switch D are assigned to the Red VLAN. If a broadcast is sent from the host connected to Switch A, Switch A floods the broadcast and every switch in the network receives it, even though Switches C, E, and F have no ports in the Red VLAN.

Figure 18-1 Flooding Traffic without VTP Pruning

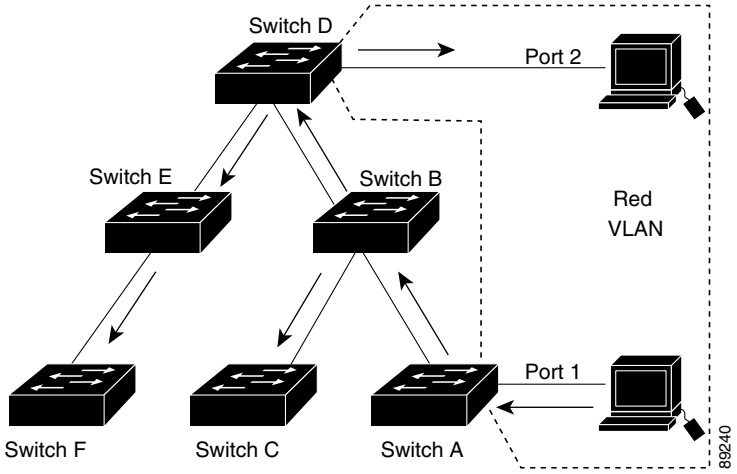

Figure 18-2 shows a switched network with VTP pruning enabled. The broadcast traffic from Switch A is not forwarded to Switches C, E, and F because traffic for the Red VLAN has been pruned on the links shown (Port 5 on Switch B and Port 4 on Switch D).

Figure 18-2 Optimized Flooded Traffic with VTP Pruning

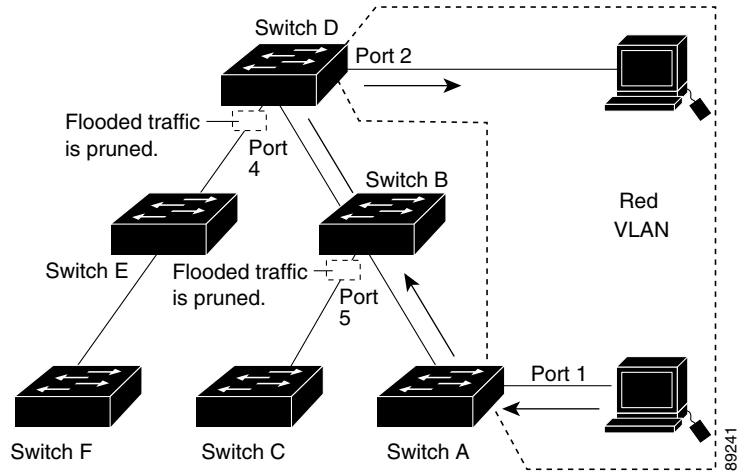

With VTP versions 1 and 2, enabling VTP pruning on a VTP server enables pruning for the entire management domain. Making VLANs pruning-eligible or pruning-ineligible affects pruning eligibility for those VLANs on that trunk only (not on all switches in the VTP domain). In VTP version 3, you must manually enable pruning on each switch in the domain.

See the "Enabling VTP Pruning" section on page 18-13. VTP pruning takes effect several seconds after you enable it. VTP pruning does not prune traffic from VLANs that are pruning-ineligible. VLAN 1 and VLANs 1002 to 1005 are always pruning-ineligible; traffic from these VLANs cannot be pruned. Extended-range VLANs (VLAN IDs higher than 1005) are also pruning-ineligible.

VTP pruning is not designed to function in VTP transparent mode. If one or more switches in the network are in VTP transparent mode, you should do one of these:

- Turn off VTP pruning in the entire network.
- Turn off VTP pruning by making all VLANs on the trunk of the switch upstream to the VTP transparent switch pruning ineligible.

To configure VTP pruning on an interface, use the **switchport trunk pruning vlan** interface configuration command. VTP pruning operates when an interface is trunking. You can set VLAN pruning-eligibility, whether or not VTP pruning is enabled for the VTP domain, whether or not any given VLAN exists, and whether or not the interface is currently trunking.

Default VTP Settings

Table 18-2 Default VTP Settings

Feature	Default Setting
VTP domain name	Null.
VTP mode (VTP version 1 and version 2)	Server.
VTP mode (VTP version 3)	The mode is the same as the mode in VTP version 1 or 2 before conversion to version 3.
VTP version	Version 1.
MST database mode	Transparent.
VTP version 3 server type	Secondary.
VTP password	None.
VTP pruning	Disabled.

VTP Configuration Guidelines

You use the **vtp** global configuration command to set the VTP password, the version, the VTP filename, the interface providing updated VTP information, the domain name, and the mode, and to disable or enable pruning. For more information about available keywords, see the command descriptions in the command reference for this release. The VTP information is saved in the VTP VLAN database. When VTP mode is transparent, the VTP domain name and mode are also saved in the switch running configuration file, and you can save it in the switch startup configuration file by entering the **copy running-config startup-config** privileged EXEC command. You must use this command if you want to save VTP mode as transparent if the switch resets.

When you save VTP information in the switch startup configuration file and restart the switch, the configuration is selected as follows:

- If the VTP mode is transparent in both the startup configuration and the VLAN database and the VTP domain name from the VLAN database matches that in the startup configuration file, the VLAN database is ignored (cleared). The VTP and VLAN configurations in the startup configuration file are used. The VLAN database revision number remains unchanged in the VLAN database.

- If the VTP mode or the domain name in the startup configuration do not match the VLAN database, the domain name and the VTP mode and configuration for the first 1005 VLANs use the VLAN database information.

Domain Names

When configuring VTP for the first time, you must always assign a domain name. You must configure all switches in the VTP domain with the same domain name. Switches in VTP transparent mode do not exchange VTP messages with other switches, and you do not need to configure a VTP domain name for them.

Note If NVRAM and DRAM storage is sufficient, all switches in a VTP domain should be in VTP server mode.

Caution Do not configure a VTP domain if all switches are operating in VTP client mode. If you configure the domain, it is impossible to make changes to the VLAN configuration of that domain. Make sure that you configure at least one switch in the VTP domain for VTP server mode.

Passwords

You can configure a password for the VTP domain, but it is not required. If you do configure a domain password, all domain switches must share the same password and you must configure the password on each switch in the management domain. Switches without a password or with the wrong password reject VTP advertisements.

If you configure a VTP password for a domain, a switch that is booted without a VTP configuration does not accept VTP advertisements until you configure it with the correct password. After the configuration, the switch accepts the next VTP advertisement that uses the same password and domain name in the advertisement.

If you are adding a new switch to an existing network with VTP capability, the new switch learns the domain name only after the applicable password has been configured on it.

Caution When you configure a VTP domain password, the management domain does not function properly if you do not assign a management domain password to each switch in the domain.

Adding a VTP Client Switch to a VTP Domain

Before adding a VTP client to a VTP domain, always verify that its VTP configuration revision number is *lower* than the configuration revision number of the other switches in the VTP domain. Switches in a VTP domain always use the VLAN configuration of the switch with the highest VTP configuration revision number. With VTP versions 1 and 2, adding a switch that has a revision number higher than the revision number in the VTP domain can erase all VLAN information from the VTP server and VTP domain. With VTP version 3, the VLAN information is not erased.

How to Configure VTP

Configuring VTP Domain and Parameters

Before You Begin

You should configure the VTP domain before configuring other VTP parameters.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**vtp domain** *domain-name*	Configures the VTP administrative-domain name. The name can be 1 to 32 characters. All switches operating in VTP server or client mode under the same administrative responsibility must be configured with the same domain name. This command is optional for modes other than server mode. VTP server mode requires a domain name. If the switch has a trunk connection to a VTP domain, the switch learns the domain name from the VTP server in the domain.
Step 3	**vtp mode** {**client** | **server** | **transparent** | **off**} {**vlan** | **mst** | **unknown**}	Configures the switch for VTP mode (client, server, transparent, or off). (Optional) Database parameters: • **vlan**—The VLAN database is the default if none are configured. • **mst**—The multiple spanning tree (MST) database. • **unknown**—An unknown database type.
Step 4	**vtp password** *password*	(Optional) Sets the password for the VTP domain. The password can be 8 to 64 characters. If you configure a VTP password, the VTP domain does not function properly if you do not assign the same password to each switch in the domain. See the "Configuring a VTP Version 3 Password" section on page 18-12 for options available with VTP version 3.
Step 1	**vtp primary-server** [**vlan** | **mst**] [**force**]	(Optional) Changes the operational state of a switch from a secondary server (the default) to a primary server and advertise the configuration to the domain. If the switch password is configured as **hidden**, you are prompted to reenter the password. • **vlan**—Selects the VLAN database as the takeover feature. This is the default. • **mst**—Selects the multiple spanning tree (MST) database as the takeover feature. • **force**—Overwrites the configuration of any conflicting servers. If you do not enter **force**, you are prompted for confirmation before the takeover.
Step 2	**end**	Returns to privileged EXEC mode.

	Command	Purpose
Step 3	**show vtp status**	Verifies your entries in the *VTP Operating Mode* and the *VTP Domain Name* fields of the display.
Step 4	**copy running-config startup-config**	(Optional) Saves the configuration in the startup configuration file. **Note** Only VTP mode and domain name are saved in the switch running configuration and can be copied to the startup configuration file.

Configuring a VTP Version 3 Password

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**vtp password** *password* [**hidden** | **secret**]	(Optional) Sets the password for the VTP domain. The password can be 8 to 64 characters. • (Optional) **hidden**—Ensures that the secret key generated from the password string is saved in the nvam:vlan.dat file. If you configure a takeover by configuring a VTP primary server, you are prompted to reenter the password. • (Optional) **secret**—Directly configures the password. The secret password must contain 32 hexadecimal characters.
Step 3	**end**	Returns to privileged EXEC mode.
Step 4	**show vtp password**	Verifies your entries.

Enabling the VTP Version

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**vtp version** {**1** | **2** | **3**}	Enables the VTP version on the switch. The default is VTP version 1.
Step 3	**end**	Returns to privileged EXEC mode.
Step 4	**show vtp status**	Verifies that the configured VTP version is enabled.
Step 5	**copy running-config startup-config**	(Optional) Saves the configuration in the startup configuration file.

Enabling VTP Pruning

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**vtp pruning**	Enables pruning in the VTP administrative domain. By default, pruning is disabled. You need to enable pruning on only one switch in VTP server mode.
Step 3	**end**	Returns to privileged EXEC mode.
Step 4	**show vtp status**	Verifies your entries in the *VTP Pruning Mode* field of the display.

Configuring VTP on a Per-Port Basis

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Identifies an interface, and enters interface configuration mode.
Step 3	**vtp**	Enables VTP on the specified port.
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**show running-config interface** *interface-id*	Verifies the change to the port.
Step 6	**show vtp status**	Verifies the configuration.

Adding a VTP Client Switch to a VTP Domain

Before You Begin

Before adding a VTP client to a VTP domain, always verify that its VTP configuration revision number is *lower* than the configuration revision number of the other switches in the VTP domain. Switches in a VTP domain always use the VLAN configuration of the switch with the highest VTP configuration revision number. With VTP versions 1 and 2, adding a switch that has a revision number higher than the revision number in the VTP domain can erase all VLAN information from the VTP server and VTP domain. With VTP version 3, the VLAN information is not erased.

	Command	Purpose
Step 1	**show vtp status**	Checks the VTP configuration revision number. If the number is 0, add the switch to the VTP domain. If the number is greater than 0, follow these steps: **a.** Write down the domain name. **b.** Write down the configuration revision number. **c.** Continue with the next steps to reset the switch configuration revision number.
Step 2	**configure terminal**	Enters global configuration mode.

	Command	Purpose
Step 3	**vtp domain** *domain-name*	Changes the domain name from the original one displayed in Step 1 to a new name.
Step 4	**end**	Updates VLAN information on the switch and resets configuration revision number to 0.
Step 5	**show vtp status**	Verifies that the configuration revision number has been reset to 0.
Step 6	**configure terminal**	Enters global configuration mode.
Step 7	**vtp domain** *domain-name*	Enters the original domain name on the switch.
Step 8	**end**	Returns to privileged EXEC mode.
Step 9	**show vtp status**	(Optional) Verifies that the domain name is the same as in Step 1 and that the configuration revision number is 0.
Step 10	After resetting the configuration revision number, add the switch to the VTP domain.	

Monitoring and Maintaining VTP

Command	Purpose
show vtp counters	Displays counters about VTP messages that have been sent and received.
show vtp devices [**conflict**]	Displays information about all VTP version 3 devices in the domain. Conflicts are VTP version 3 devices with conflicting primary servers. The **show vtp devices** command does not display information when the switch is in transparent or off mode.
show vtp interface [*interface-id*]	Displays VTP status and configuration for all interfaces or the specified interface.
show vtp password	Displays the VTP password. The form of the password displayed depends on whether or not the **hidden** keyword was entered and if encryption is enabled on the switch.
show vtp status	Displays the VTP switch configuration information.

Configuration Examples for Configuring VTP

Configuring a VTP Server: Example

This example shows how to configure the switch as a VTP server with the domain name *eng_group* and the password *mypassword*:

```
Switch(config)# vtp domain eng_group
Setting VTP domain name to eng_group.
Switch(config)# vtp mode server
Setting device to VTP Server mode for VLANS.
```

```
Switch(config)# vtp password mypassword
Setting device VLAN database password to mypassword.
Switch(config)# end
```

Configuring a Hidden VTP Password: Example

This example shows how to configure a hidden password and how it appears:

```
Switch(config)# vtp password mypassword hidden
Generating the secret associated to the password.
Switch(config)# end
Switch# show vtp password
VTP password: 89914640C8D90868B6A0D8103847A733
```

Configuring a VTP Version 3 Primary Server: Example

This example shows how to configure a switch as the primary server for the VLAN database (the default) when a hidden or secret password was configured:

```
Switch# vtp primary vlan
Enter VTP password: mypassword
This switch is becoming Primary server for vlan feature in the VTP  domain

VTP Database Conf Switch ID      Primary Server Revision System Name
------------ ---- --------------- -------------- -------- --------------------
VLANDB       Yes  00d0.00b8.1400=00d0.00b8.1400 1        stp7

Do you want to continue (y/n) [n]? y
```

Additional References for Configuring VTP

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
VLAN configuration	"Configuring VLANs"

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

CHAPTER 19

Configuring Voice VLAN

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Information About Configuring Voice VLAN

Voice VLAN

The voice VLAN feature enables access ports to carry IP voice traffic from an IP phone. When the switch is connected to a Cisco 7960 IP Phone, the phone sends voice traffic with Layer 3 IP precedence and Layer 2 class of service (CoS) values, which are both set to 5 by default. Because the sound quality of a Cisco IP phone call can deteriorate if the data is unevenly sent, the switch supports quality of service (QoS) based on IEEE 802.1p CoS. QoS uses classification and scheduling to send network traffic from the switch in a predictable manner. Voice VLAN is referred to as an *auxiliary VLAN* in some Catalyst 6500 family switch documentation.

The Cisco 7960 IP Phone is a configurable device, and you can configure it to forward traffic with an IEEE 802.1p priority. You can configure the switch to trust or override the traffic priority assigned by a Cisco IP phone.

The Cisco IP phone contains an integrated three-port 10/100 switch as shown in Figure 19-1. The ports provide dedicated connections to these devices:

- Port 1 connects to the switch or other voice-over-IP (VoIP) device.
- Port 2 is an internal 10/100 interface that carries the IP phone traffic.
- Port 3 (access port) connects to a PC or other device.

Figure 19-1 Cisco 7960 IP Phone Connected to a Switch

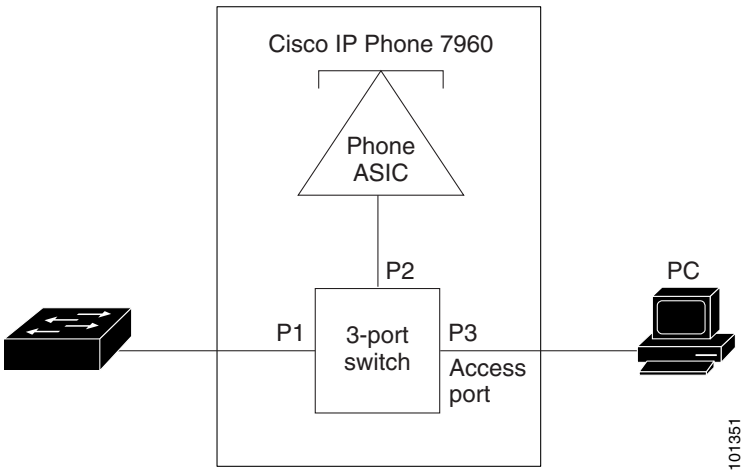

Cisco IP Phone Voice Traffic

You can configure an access port with an attached Cisco IP phone to use one VLAN for voice traffic and another VLAN for data traffic from a device attached to the phone. You can configure access ports on the switch to send Cisco Discovery Protocol (CDP) packets that instruct an attached phone to send voice traffic to the switch in any of these ways:

- In the voice VLAN tagged with a Layer 2 CoS priority value
- In the access VLAN tagged with a Layer 2 CoS priority value
- In the access VLAN, untagged (no Layer 2 CoS priority value)

Note In all configurations, the voice traffic carries a Layer 3 IP precedence value (the default is 5 for voice traffic and 3 for voice control traffic).

You can configure a port connected to the Cisco IP phone to send CDP packets to the phone to configure the way in which the phone sends voice traffic. The phone can carry voice traffic in IEEE 802.1Q frames for a specified voice VLAN with a Layer 2 CoS value. It can use IEEE 802.1p priority tagging to give voice traffic a higher priority and forward all voice traffic through the native (access) VLAN. The Cisco IP phone can also send untagged voice traffic or use its own configuration to send voice traffic in the access VLAN. In all configurations, the voice traffic carries a Layer 3 IP precedence value (the default is 5).

Cisco IP Phone Data Traffic

The switch can also process tagged data traffic (traffic in IEEE 802.1Q or IEEE 802.1p frame types) from the device attached to the access port on the Cisco IP phone (see Figure 19-1). You can configure Layer 2 access ports on the switch to send CDP packets that instruct the attached phone to configure the phone access port in one of these modes:

- In trusted mode, all traffic received through the access port on the Cisco IP phone passes through the phone unchanged.
- In untrusted mode, all traffic in IEEE 802.1Q or IEEE 802.1p frames received through the access port on the Cisco IP phone receive a configured Layer 2 CoS value. The default Layer 2 CoS value is 0. Untrusted mode is the default.

Note Untagged traffic from the device attached to the Cisco IP phone passes through the phone unchanged, regardless of the trust state of the access port on the phone.

Default Voice VLAN Configuration

The voice VLAN feature is disabled by default.

When the voice VLAN feature is enabled, all untagged traffic is sent according to the default CoS priority of the port.

The CoS value is not trusted for IEEE 802.1p or IEEE 802.1Q tagged traffic.

Voice VLAN Configuration Guidelines

- Voice VLAN configuration is only supported on switch access ports; voice VLAN configuration is not supported on trunk ports.

Note Trunk ports can carry any number of voice VLANs, similar to regular VLANs. The configuration of voice VLANs is not required on trunk ports.

- The voice VLAN should be present and active on the switch for the IP phone to correctly communicate on the voice VLAN. Use the **show vlan** privileged EXEC command to see if the VLAN is present (listed in the display). If the VLAN is not listed, see Chapter 17, "Configuring VLANs," for information on how to create the voice VLAN.
- Before you enable voice VLAN, we recommend that you enable QoS on the switch by entering the **mls qos** global configuration command and configure the port trust state to trust by entering the **mls qos trust cos** interface configuration command. If you use the auto-QoS feature, these settings are automatically configured. For more information, see Chapter 38, "Configuring Standard QoS."
- You must enable CDP on the switch port connected to the Cisco IP phone to send the configuration to the phone. (CDP is globally enabled by default on all switch interfaces.)
- The Port Fast feature is automatically enabled when voice VLAN is configured. When you disable voice VLAN, the Port Fast feature is not automatically disabled.
- If the Cisco IP phone and a device attached to the phone are in the same VLAN, they must be in the same IP subnet. These conditions indicate that they are in the same VLAN:

  - They both use IEEE 802.1p or untagged frames.
  - The Cisco IP phone uses IEEE 802.1p frames, and the device uses untagged frames.
  - The Cisco IP phone uses untagged frames, and the device uses IEEE 802.1p frames.
  - The Cisco IP phone uses IEEE 802.1Q frames, and the voice VLAN is the same as the access VLAN.
- The Cisco IP phone and a device attached to the phone cannot communicate if they are in the same VLAN and subnet but use different frame types because traffic in the same subnet is not routed (routing would eliminate the frame type difference).
- You cannot configure static secure MAC addresses in the voice VLAN.
- Voice VLAN ports can also be these port types:
  - Dynamic access port. See the "Configuring Dynamic-Access Ports on VMPS Clients" section on page 17-23 for more information.
  - IEEE 802.1x authenticated port. See the "Configuring 802.1x Readiness Check" section on page 13-36 for more information.

    Note If you enable IEEE 802.1x on an access port on which a voice VLAN is configured and to which a Cisco IP phone is connected, the phone loses connectivity to the switch for up to 30 seconds.

  - Protected port. See the "Configuring Protected Ports" section on page 29-10 for more information.
  - A source or destination port for a SPAN or RSPAN session.
  - Secure port. See the "Configuring Port Security" section on page 29-11 for more information.

    Note When you enable port security on an interface that is also configured with a voice VLAN, you must set the maximum allowed secure addresses on the port to two plus the maximum number of secure addresses allowed on the access VLAN. When the port is connected to a Cisco IP phone, the phone requires up to two MAC addresses. The phone address is learned on the voice VLAN and might also be learned on the access VLAN. Connecting a PC to the phone requires additional MAC addresses.

Port Connection to a Cisco 7960 IP Phone

Because a Cisco 7960 IP Phone also supports a connection to a PC or other device, a port connecting the switch to a Cisco IP phone can carry mixed traffic. You can configure a port to decide how the Cisco IP phone carries voice traffic and data traffic.

Priority of Incoming Data Frames

You can connect a PC or other data device to a Cisco IP phone port. To process tagged data traffic (in IEEE 802.1Q or IEEE 802.1p frames), you can configure the switch to send CDP packets to instruct the phone how to send data packets from the device attached to the access port on the Cisco IP phone. The PC can generate packets with an assigned CoS value. You can configure the phone to not change (trust) or to override (not trust) the priority of frames arriving on the phone port from connected devices.

How to Configure VTP

Configuring Cisco IP Phone for Voice Traffic

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface connected to the phone, and enters interface configuration mode.
Step 3	**mls qos trust cos**	Configures the interface to classify incoming traffic packets by using the packet CoS value. For untagged packets, the port default CoS value is used. **Note** Before configuring the port trust state, you must first globally enable QoS by using the **mls qos** global configuration command.
Step 4	**switchport voice vlan** {*vlan-id* | **dot1p** | **none** | **untagged**}}	Configures how the Cisco IP phone carries voice traffic: • *vlan-id*—Configures the phone to forward all voice traffic through the specified VLAN. By default, the Cisco IP phone forwards the voice traffic with an IEEE 802.1Q priority of 5. Valid VLAN IDs are 1 to 4096. • **dot1p**—Configures the phone to use IEEE 802.1p priority tagging for voice traffic and to use the default native VLAN (VLAN 0) to carry all traffic. By default, the Cisco IP phone forwards the voice traffic with an IEEE 802.1p priority of 5. • **none**—Allows the phone to use its own configuration to send untagged voice traffic. • **untagged**—Configures the phone to send untagged voice traffic.
Step 5	**end**	Returns to privileged EXEC mode.

Configuring the Priority of Incoming Data Frames

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface connected to the Cisco IP phone, and enters interface configuration mode.
Step 3	**switchport priority extend** {**cos** *value* | **trust**}	Sets the priority of data traffic received from the Cisco IP phone access port: • **cos** *value*—Configures the phone to override the priority received from the PC or the attached device with the specified CoS value. The value is a number from 0 to 7, with 7 as the highest priority. The default priority is **cos** 0. • **trust**—Configures the phone access port to trust the priority received from the PC or the attached device.
Step 4	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining Voice VLAN

Command	Purpose
show interfaces *interface-id* **switchport**	Verifies your entries.
copy running-config startup-config	Saves your entries in the configuration file.

Configuration Examples for Configuring Voice VLAN

Configuring a Cisco IP Phone for Voice Traffic: Example

This example shows how to configure a port connected to a Cisco IP phone to use the CoS value to classify incoming traffic, to use IEEE 802.1p priority tagging for voice traffic, and to use the default native VLAN (VLAN 0) to carry all traffic:

```
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# mls qos trust cos
Switch(config-if)# switchport voice vlan dot1p
Switch(config-if)# end
```

Configuring the Cisco IP Phone Priority of Incoming Data Frames: Example

This example shows how to configure a port connected to a Cisco IP phone to not change the priority of frames received from the PC or the attached device:

```
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# switchport priority extend trust
Switch(config-if)# end
```

Additional References for Configuring Voice VLAN

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
QoS configuration	*Chapter 38, "Configuring Standard QoS"*
VLAN configuration	*Chapter 17, "Configuring VLANs"*
Dynamic access port configuration	"Configuring Dynamic-Access Ports on VMPS Clients" section on page 17-23
IEEE 802.1x authenticated port configuration	"Configuring 802.1x Readiness Check" section on page 13-36

Related Topic	Document Title
Protected port configuration	"Configuring Protected Ports" section on page 29-10
Secure port configuration	"Configuring Port Security" section on page 29-11

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

CHAPTER 20

Configuring STP

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for Configuring STP

When you configure VTP, you must configure a trunk port so that the switch can send and receive VTP advertisements to and from other switches in the domain.

For more information, see the "Configuring an Ethernet Interface as a Trunk Port" section on page 17-19.

Restrictions for Configuring STP

- If you are configuring VTP on a cluster member switch to a VLAN, use the **rcommand** privileged EXEC command to log in to the member switch.
- In VTP versions 1 and 2, when you configure extended-range VLANs on the switch, the switch must be in VTP transparent mode. VTP version 3 also supports creating extended-range VLANs in client or server mode.

Information About Configuring STP

This chapter describes how to configure the Spanning Tree Protocol (STP) on port-based VLANs on the switch. The switch can use either the per-VLAN spanning-tree plus (PVST+) protocol based on the IEEE 802.1D standard and Cisco proprietary extensions, or the rapid per-VLAN spanning-tree plus (rapid-PVST+) protocol based on the IEEE 802.1w standard.

STP

STP is a Layer 2 link management protocol that provides path redundancy while preventing loops in the network. For a Layer 2 Ethernet network to function properly, only one active path can exist between any two stations. Multiple active paths among end stations cause loops in the network. If a loop exists in the network, end stations might receive duplicate messages. Switches might also learn end-station MAC addresses on multiple Layer 2 interfaces. These conditions result in an unstable network. Spanning-tree operation is transparent to end stations, which cannot detect whether they are connected to a single LAN segment or a switched LAN of multiple segments.

The STP uses a spanning-tree algorithm to select one switch of a redundantly connected network as the root of the spanning tree. The algorithm calculates the best loop-free path through a switched Layer 2 network by assigning a role to each port based on the role of the port in the active topology:

- Root—A forwarding port elected for the spanning-tree topology
- Designated—A forwarding port elected for every switched LAN segment
- Alternate—A blocked port providing an alternate path to the root bridge in the spanning tree
- Backup—A blocked port in a loopback configuration

The switch that has *all* of its ports as the designated role or as the backup role is the root switch. The switch that has at least *one* of its ports in the designated role is called the designated switch.

Spanning tree forces redundant data paths into a standby (blocked) state. If a network segment in the spanning tree fails and a redundant path exists, the spanning-tree algorithm recalculates the spanning-tree topology and activates the standby path. Switches send and receive spanning-tree frames, called bridge protocol data units (BPDUs), at regular intervals. The switches do not forward these frames but use them to construct a loop-free path. BPDUs contain information about the sending switch and its ports, including switch and MAC addresses, switch priority, port priority, and path cost. Spanning tree uses this information to elect the root switch and root port for the switched network and the root port and designated port for each switched segment.

When two ports on a switch are part of a loop, the spanning-tree port priority and path cost settings control which port is put in the forwarding state and which is put in the blocking state. The spanning-tree port priority value represents the location of a port in the network topology and how well it is located to pass traffic. The path cost value represents the media speed.

Note The default is for the switch to send keepalive messages (to ensure the connection is up) only on interfaces that do not have small form-factor pluggable (SFP) modules. You can use the [**no**] **keepalive** interface configuration command to change the default for an interface.

Spanning-Tree Topology and BPDUs

The stable, active spanning-tree topology of a switched network is controlled by these elements:

- The unique bridge ID (switch priority and MAC address) associated with each VLAN on each switch.
- The spanning-tree path cost to the root switch.
- The port identifier (port priority and MAC address) associated with each Layer 2 interface.

When the switches in a network are powered up, each functions as the root switch. Each switch sends a configuration BPDU through all of its ports. The BPDUs communicate and compute the spanning-tree topology. Each configuration BPDU contains this information:

- The unique bridge ID of the switch that the sending switch identifies as the root switch
- The spanning-tree path cost to the root
- The bridge ID of the sending switch
- Message age
- The identifier of the sending interface
- Values for the hello, forward delay, and max-age protocol timers

When a switch receives a configuration BPDU that contains *superior* information (lower bridge ID, lower path cost, and so forth), it stores the information for that port. If this BPDU is received on the root port of the switch, the switch also forwards it with an updated message to all attached LANs for which it is the designated switch.

If a switch receives a configuration BPDU that contains *inferior* information to that currently stored for that port, it discards the BPDU. If the switch is a designated switch for the LAN from which the inferior BPDU was received, it sends that LAN a BPDU containing the up-to-date information stored for that port. In this way, inferior information is discarded, and superior information is propagated on the network.

A BPDU exchange results in these actions:

- One switch in the network is elected as the root switch (the logical center of the spanning-tree topology in a switched network).

  For each VLAN, the switch with the highest switch priority (the lowest numerical priority value) is elected as the root switch. If all switches are configured with the default priority (32768), the switch with the lowest MAC address in the VLAN becomes the root switch. The switch priority value occupies the most significant bits of the bridge ID, as shown in Table 20-1 on page 20-4.
- A root port is selected for each switch (except the root switch). This port provides the best path (lowest cost) when the switch forwards packets to the root switch.
- The shortest distance to the root switch is calculated for each switch based on the path cost.
- A designated switch for each LAN segment is selected. The designated switch incurs the lowest path cost when forwarding packets from that LAN to the root switch. The port through which the designated switch is attached to the LAN is called the designated port.

All paths that are not needed to reach the root switch from anywhere in the switched network are placed in the spanning-tree blocking mode.

Bridge ID, Switch Priority, and Extended System ID

The IEEE 802.1D standard requires that each switch has an unique bridge identifier (bridge ID), which controls the selection of the root switch. Because each VLAN is considered as a different *logical bridge* with PVST+ and rapid PVST+, the same switch must have a different bridge IDs for each configured VLAN. Each VLAN on the switch has a unique 8-byte bridge ID. The 2 most-significant bytes are used for the switch priority, and the remaining 6 bytes are derived from the switch MAC address.

The switch supports the IEEE 802.1t spanning-tree extensions, and some of the bits previously used for the switch priority are now used as the VLAN identifier. The result is that fewer MAC addresses are reserved for the switch, and a larger range of VLAN IDs can be supported, all while maintaining the uniqueness of the bridge ID. As shown in Table 20-1, the 2 bytes previously used for the switch priority are reallocated into a 4-bit priority value and a 12-bit extended system ID value equal to the VLAN ID.

Table 20-1 Switch Priority Value and Extended System ID

Switch Priority Value				Extended System ID (Set Equal to the VLAN ID)											
Bit 16	Bit 15	Bit 14	Bit 13	Bit 12	Bit 11	Bit 10	Bit 9	Bit 8	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1
32768	16384	8192	4096	2048	1024	512	256	128	64	32	16	8	4	2	1

Spanning tree uses the extended system ID, the switch priority, and the allocated spanning-tree MAC address to make the bridge ID unique for each VLAN.

Support for the extended system ID affects how you manually configure the root switch, the secondary root switch, and the switch priority of a VLAN. For example, when you change the switch priority value, you change the probability that the switch will be elected as the root switch. Configuring a higher value decreases the probability; a lower value increases the probability. For more information, see the "Configuring the Root Switch" section on page 20-15, the "Configuring a Secondary Root Switch" section on page 20-16, and the "Configuring Optional STP Parameters" section on page 20-17.

Spanning-Tree Interface States

Propagation delays can occur when protocol information passes through a switched LAN. As a result, topology changes can take place at different times and at different places in a switched network. When an interface transitions directly from nonparticipation in the spanning-tree topology to the forwarding state, it can create temporary data loops. Interfaces must wait for new topology information to propagate through the switched LAN before starting to forward frames. They must allow the frame lifetime to expire for forwarded frames that have used the old topology.

Each Layer 2 interface on a switch using spanning tree exists in one of these states:

- Blocking—The interface does not participate in frame forwarding.
- Listening—The first transitional state after the blocking state when the spanning tree decides that the interface should participate in frame forwarding.
- Learning—The interface prepares to participate in frame forwarding.
- Forwarding—The interface forwards frames.
- Disabled—The interface is not participating in spanning tree because of a shutdown port, no link on the port, or no spanning-tree instance running on the port.

An interface moves through these states:

- From initialization to blocking
- From blocking to listening or to disabled
- From listening to learning or to disabled
- From learning to forwarding or to disabled
- From forwarding to disabled

Figure 20-1 illustrates how an interface moves through the states.

Figure 20-1 Spanning-Tree Interface States

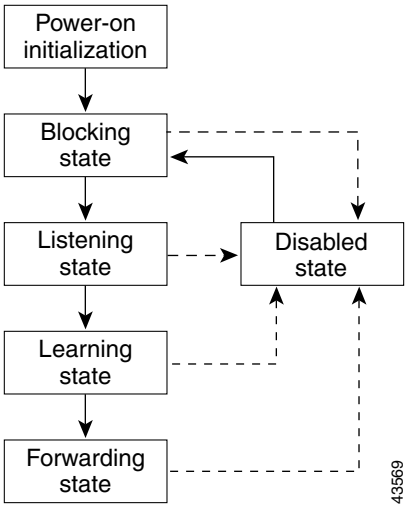

When you power up the switch, spanning tree is enabled by default, and every interface in the switch, VLAN, or network goes through the blocking state and the transitory states of listening and learning. Spanning tree stabilizes each interface at the forwarding or blocking state.

When the spanning-tree algorithm places a Layer 2 interface in the forwarding state, this process occurs:

1. The interface is in the listening state while spanning tree waits for protocol information to move the interface to the blocking state.
2. While spanning tree waits the forward-delay timer to expire, it moves the interface to the learning state and resets the forward-delay timer.
3. In the learning state, the interface continues to block frame forwarding as the switch learns end-station location information for the forwarding database.
4. When the forward-delay timer expires, spanning tree moves the interface to the forwarding state, where both learning and frame forwarding are enabled.

Blocking State

A Layer 2 interface in the blocking state does not participate in frame forwarding. After initialization, a BPDU is sent to each switch interface. A switch initially functions as the root until it exchanges BPDUs with other switches. This exchange establishes which switch in the network is the root or root switch. If there is only one switch in the network, no exchange occurs, the forward-delay timer expires, and the interface moves to the listening state. An interface always enters the blocking state after switch initialization.

An interface in the blocking state performs these functions:

- Discards frames received on the interface
- Discards frames switched from another interface for forwarding
- Does not learn addresses
- Receives BPDUs

Listening State

The listening state is the first state a Layer 2 interface enters after the blocking state. The interface enters this state when the spanning tree decides that the interface should participate in frame forwarding.

An interface in the listening state performs these functions:

- Discards frames received on the interface
- Discards frames switched from another interface for forwarding
- Does not learn addresses
- Receives BPDUs

Learning State

A Layer 2 interface in the learning state prepares to participate in frame forwarding. The interface enters the learning state from the listening state.

An interface in the learning state performs these functions:

- Discards frames received on the interface
- Discards frames switched from another interface for forwarding
- Learns addresses
- Receives BPDUs

Forwarding State

A Layer 2 interface in the forwarding state forwards frames. The interface enters the forwarding state from the learning state.

An interface in the forwarding state performs these functions:

- Receives and forwards frames received on the interface
- Forwards frames switched from another interface
- Learns addresses
- Receives BPDUs

Disabled State

A Layer 2 interface in the disabled state does not participate in frame forwarding or in the spanning tree. An interface in the disabled state is nonoperational.

A disabled interface performs these functions:

- Discards frames received on the interface
- Discards frames switched from another interface for forwarding
- Does not learn addresses
- Does not receive BPDUs

How a Switch or Port Becomes the Root Switch or Root Port

If all switches in a network are enabled with default spanning-tree settings, the switch with the lowest MAC address becomes the root switch. In Figure 20-2, Switch A is elected as the root switch because the switch priority of all the switches is set to the default (32768) and Switch A has the lowest MAC address. However, because of traffic patterns, number of forwarding interfaces, or link types, Switch A might not be the ideal root switch. By increasing the priority (lowering the numerical value) of the ideal switch so that it becomes the root switch, you force a spanning-tree recalculation to form a new topology with the ideal switch as the root.

Figure 20-2 Spanning-Tree Topology

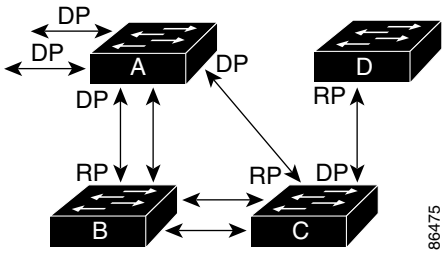

RP = Root Port
DP = Designated Port

When the spanning-tree topology is calculated based on default parameters, the path between source and destination end stations in a switched network might not be ideal. For instance, connecting higher-speed links to an interface that has a higher number than the root port can cause a root-port change. The goal is to make the fastest link the root port.

For example, assume that one port on Switch B is a Gigabit Ethernet link and that another port on Switch B (a 10/100 link) is the root port. Network traffic might be more efficient over the Gigabit Ethernet link. By changing the spanning-tree port priority on the Gigabit Ethernet port to a higher priority (lower numerical value) than the root port, the Gigabit Ethernet port becomes the new root port.

Spanning Tree and Redundant Connectivity

You can create a redundant backbone with spanning tree by connecting two switch interfaces to another device or to two different devices, as shown in Figure 20-3. Spanning tree automatically disables one interface but enables it if the other one fails. If one link is high-speed and the other is low-speed, the low-speed link is always disabled. If the speeds are the same, the port priority and port ID are added together, and spanning tree disables the link with the lowest value.

Figure 20-3 Spanning Tree and Redundant Connectivity

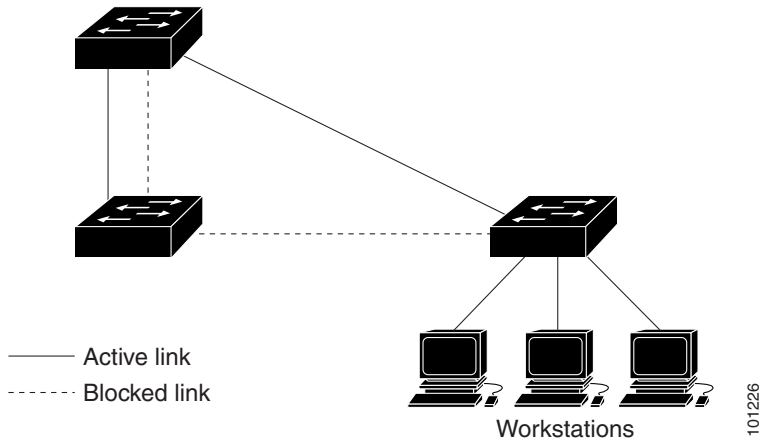

You can also create redundant links between switches by using EtherChannel groups. For more information, see Chapter 40, "Configuring EtherChannels."

Spanning-Tree Address Management

IEEE 802.1D specifies 17 multicast addresses, ranging from 0x00180C2000000 to 0x0180C2000010, to be used by different bridge protocols. These addresses are static addresses that cannot be removed.

Regardless of the spanning-tree state, each switch receives but does not forward packets destined for addresses between 0x0180C2000000 and 0x0180C200000F.

If spanning tree is enabled, the CPU on the switch receives packets destined for 0x0180C2000000 and 0x0180C2000010. If spanning tree is disabled, the switch forwards those packets as unknown multicast addresses.

Accelerated Aging to Retain Connectivity

The default for aging dynamic addresses is 5 minutes, the default setting of the **mac address-table aging-time** global configuration command. However, a spanning-tree reconfiguration can cause many station locations to change. Because these stations could be unreachable for 5 minutes or more during a reconfiguration, the address-aging time is accelerated so that station addresses can be dropped from the address table and then relearned. The accelerated aging is the same as the forward-delay parameter value (**spanning-tree vlan** *vlan-id* **forward-time** *seconds* global configuration command) when the spanning tree reconfigures.

Because each VLAN is a separate spanning-tree instance, the switch accelerates aging on a per-VLAN basis. A spanning-tree reconfiguration on one VLAN can cause the dynamic addresses learned on that VLAN to be subject to accelerated aging. Dynamic addresses on other VLANs can be unaffected and remain subject to the aging interval entered for the switch.

Spanning-Tree Modes and Protocols

The switch supports these spanning-tree modes and protocols:

- PVST+—This spanning-tree mode is based on the IEEE 802.1D standard and Cisco proprietary extensions. It is the default spanning-tree mode used on all Ethernet port-based VLANs. The PVST+ runs on each VLAN on the switch up to the maximum supported, ensuring that each has a loop-free path through the network.

  The PVST+ provides Layer 2 load balancing for the VLAN on which it runs. You can create different logical topologies by using the VLANs on your network to ensure that all of your links are used but that no one link is oversubscribed. Each instance of PVST+ on a VLAN has a single root switch. This root switch propagates the spanning-tree information associated with that VLAN to all other switches in the network. Because each switch has the same information about the network, this process ensures that the network topology is maintained.

- Rapid PVST+—This spanning-tree mode is the same as PVST+ except that is uses a rapid convergence based on the IEEE 802.1w standard. To provide rapid convergence, the rapid PVST+ immediately deletes dynamically learned MAC address entries on a per-port basis upon receiving a topology change. By contrast, PVST+ uses a short aging time for dynamically learned MAC address entries.

  The rapid PVST+ uses the same configuration as PVST+ (except where noted), and the switch needs only minimal extra configuration. The benefit of rapid PVST+ is that you can migrate a large PVST+ install base to rapid PVST+ without having to learn the complexities of the MSTP configuration and without having to reprovision your network. In rapid-PVST+ mode, each VLAN runs its own spanning-tree instance up to the maximum supported.

- MSTP—This spanning-tree mode is based on the IEEE 802.1s standard. You can map multiple VLANs to the same spanning-tree instance, which reduces the number of spanning-tree instances required to support a large number of VLANs. The MSTP runs on top of the RSTP (based on IEEE 802.1w), which provides for rapid convergence of the spanning tree by eliminating the forward delay and by quickly transitioning root ports and designated ports to the forwarding state. You cannot run MSTP without RSTP.

  The most common initial deployment of MSTP is in the backbone and distribution layers of a Layer 2 switched network. For more information, see Chapter 21, “Configuring MSTP.”

For information about the number of supported spanning-tree instances, see the next section.

Supported Spanning-Tree Instances

In PVST+ or rapid-PVST+ mode, the switch supports up to 128 spanning-tree instances.

In MSTP mode, the switch supports up to 65 MST instances. The number of VLANs that can be mapped to a particular MST instance is unlimited.

For information about how spanning tree interoperates with the VLAN Trunking Protocol (VTP), see the “Changing the Spanning-Tree Mode” section on page 20-14.

Spanning-Tree Interoperability and Backward Compatibility

Table 20-2 lists the interoperability and compatibility among the supported spanning-tree modes in a network.

Table 20-2 PVST+, MSTP, and Rapid-PVST+ Interoperability

	PVST+	MSTP	Rapid PVST+
PVST+	Yes	Yes (with restrictions)	Yes (reverts to PVST+)
MSTP	Yes (with restrictions)	Yes	Yes (reverts to PVST+)
Rapid PVST+	Yes (reverts to PVST+)	Yes (reverts to PVST+)	Yes

In a mixed MSTP and PVST+ network, the common spanning-tree (CST) root must be inside the MST backbone, and a PVST+ switch cannot connect to multiple MST regions.

When a network contains switches running rapid PVST+ and switches running PVST+, we recommend that the rapid-PVST+ switches and PVST+ switches be configured for different spanning-tree instances. In the rapid-PVST+ spanning-tree instances, the root switch must be a rapid-PVST+ switch. In the PVST+ instances, the root switch must be a PVST+ switch. The PVST+ switches should be at the edge of the network.

STP and IEEE 802.1Q Trunks

The IEEE 802.1Q standard for VLAN trunks imposes some limitations on the spanning-tree strategy for a network. The standard requires only one spanning-tree instance for *all* VLANs allowed on the trunks. However, in a network of Cisco switches connected through IEEE 802.1Q trunks, the switches maintain one spanning-tree instance for *each* VLAN allowed on the trunks.

When you connect a Cisco switch to a non-Cisco device through an IEEE 802.1Q trunk, the Cisco switch uses PVST+ to provide spanning-tree interoperability. If rapid PVST+ is enabled, the switch uses it instead of PVST+. The switch combines the spanning-tree instance of the IEEE 802.1Q VLAN of the trunk with the spanning-tree instance of the non-Cisco IEEE 802.1Q switch.

However, all PVST+ or rapid-PVST+ information is maintained by Cisco switches separated by a cloud of non-Cisco IEEE 802.1Q switches. The non-Cisco IEEE 802.1Q cloud separating the Cisco switches is treated as a single trunk link between the switches.

PVST+ is automatically enabled on IEEE 802.1Q trunks, and no user configuration is required. The external spanning-tree behavior on access ports is not affected by PVST+.

VLAN-Bridge Spanning Tree

Cisco VLAN-bridge spanning tree is used with the fallback bridging feature (bridge groups), which forwards non-IP protocols such as DECnet between two or more VLAN bridge domains or routed ports. The VLAN-bridge spanning tree allows the bridge groups to form a spanning tree on top of the individual VLAN spanning trees to prevent loops from forming if there are multiple connections among VLANs. It also prevents the individual spanning trees from the VLANs being bridged from collapsing into a single spanning tree.

To support VLAN-bridge spanning tree, some of the spanning-tree timers are increased.

Default Spanning-Tree Settings

Table 20-3 Default Spanning-Tree Settings

Feature	Default Setting
Enable state	Enabled on VLAN 1.
Spanning-tree mode	PVST+. (Rapid PVST+ and MSTP are disabled.)
Switch priority	32768.
Spanning-tree port priority (configurable on a per-interface basis)	128.
Spanning-tree port cost (configurable on a per-interface basis)	1000 Mb/s: 4. 100 Mb/s: 19. 10 Mb/s: 100.
Spanning-tree VLAN port priority (configurable on a per-VLAN basis)	128.
Spanning-tree VLAN port cost (configurable on a per-VLAN basis)	1000 Mb/s: 4. 100 Mb/s: 19. 10 Mb/s: 100.
Spanning-tree timers	Hello time: 2 seconds. Forward-delay time: 15 seconds. Maximum-aging time: 20 seconds. Transmit hold count: 6 BPDUs

Disabling Spanning Tree

Spanning tree is enabled by default on VLAN 1 and on all newly created VLANs up to the spanning-tree limit specified in the "Supported Spanning-Tree Instances" section on page 20-9. Disable spanning tree only if you are sure there are no loops in the network topology.

Caution When spanning tree is disabled and loops are present in the topology, excessive traffic and indefinite packet duplication can drastically reduce network performance.

Root Switch

The switch maintains a separate spanning-tree instance for each active VLAN configured on it. A bridge ID, consisting of the switch priority and the switch MAC address, is associated with each instance. For each VLAN, the switch with the lowest bridge ID becomes the root switch for that VLAN.

To configure a switch to become the root for the specified VLAN, use the **spanning-tree vlan** *vlan-id* **root** global configuration command to modify the switch priority from the default value (32768) to a significantly lower value. When you enter this command, the software checks the switch priority of the root switches for each VLAN. Because of the extended system ID support, the switch sets its own priority for the specified VLAN to 24576 if this value will cause this switch to become the root for the specified VLAN.

If any root switch for the specified VLAN has a switch priority lower than 24576, the switch sets its own priority for the specified VLAN to 4096 less than the lowest switch priority. (4096 is the value of the least-significant bit of a 4-bit switch priority value as shown in Table 20-1 on page 20-4.)

Note The **spanning-tree vlan** *vlan-id* **root** global configuration command fails if the value necessary to be the root switch is less than 1.

Note If your network consists of switches that both do and do not support the extended system ID, it is unlikely that the switch with the extended system ID support will become the root switch. The extended system ID increases the switch priority value every time the VLAN number is greater than the priority of the connected switches running older software.

Note The root switch for each spanning-tree instance should be a backbone or distribution switch. Do not configure an access switch as the spanning-tree primary root.

Use the **diameter** keyword to specify the Layer 2 network diameter (that is, the maximum number of switch hops between any two end stations in the Layer 2 network). When you specify the network diameter, the switch automatically sets an optimal hello time, forward-delay time, and maximum-age time for a network of that diameter, which can significantly reduce the convergence time. You can use the **hello** keyword to override the automatically calculated hello time.

Note After configuring the switch as the root switch, we recommend that you avoid manually configuring the hello time, forward-delay time, and maximum-age time through the **spanning-tree vlan** *vlan-id* **hello-time**, **spanning-tree vlan** *vlan-id* **forward-time**, and the **spanning-tree vlan** *vlan-id* **max-age** global configuration commands.

Secondary Root Switch

When you configure a switch as the secondary root, the switch priority is modified from the default value (32768) to 28672. The switch is then likely to become the root switch for the specified VLAN if the primary root switch fails. This is assuming that the other network switches use the default switch priority of 32768 and therefore are unlikely to become the root switch.

You can execute this command on more than one switch to configure multiple backup root switches. Use the same network diameter and hello-time values that you used when you configured the primary root switch with the **spanning-tree vlan** *vlan-id* **root primary** global configuration command.

Port Priority

If a loop occurs, spanning tree uses the port priority when selecting an interface to put into the forwarding state. You can assign higher priority values (lower numerical values) to interfaces that you want selected first and lower priority values (higher numerical values) that you want selected last. If all interfaces have the same priority value, spanning tree puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.

Path Cost

The spanning-tree path cost default value is derived from the media speed of an interface. If a loop occurs, spanning tree uses cost when selecting an interface to put in the forwarding state. You can assign lower cost values to interfaces that you want selected first and higher cost values that you want selected last. If all interfaces have the same cost value, spanning tree puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.

Spanning-Tree Timers

Table 20-4 Spanning-Tree Timers

Variable	Description
Hello timer	Controls how often the switch broadcasts hello messages to other switches.
Forward-delay timer	Controls how long each of the listening and learning states last before the interface begins forwarding.
Maximum-age timer	Controls the amount of time the switch stores protocol information received on an interface.
Transmit hold count	Controls the number of BPDUs that can be sent before pausing for 1 second.

Spanning-Tree Configuration Guidelines

If more VLANs are defined in the VTP than there are spanning-tree instances, you can enable PVST+ or rapid PVST+ on only 128 VLANs on the switch. The remaining VLANs operate with spanning tree disabled. However, you can map multiple VLANs to the same spanning-tree instances by using MSTP. For more information, see Chapter 21, "Configuring MSTP."

If 128 instances of spanning tree are already in use, you can disable spanning tree on one of the VLANs and then enable it on the VLAN where you want it to run. Use the **no spanning-tree vlan** *vlan-id* global configuration command to disable spanning tree on a specific VLAN, and use the **spanning-tree vlan** *vlan-id* global configuration command to enable spanning tree on the desired VLAN.

Caution Switches that are not running spanning tree still forward BPDUs that they receive so that the other switches on the VLAN that have a running spanning-tree instance can break loops. Therefore, spanning tree must be running on enough switches to break all the loops in the network; for example, at least one switch on each loop in the VLAN must be running spanning tree. It is not absolutely necessary to run spanning tree on all switches in the VLAN. However, if you are running spanning tree only on a minimal set of switches, an incautious change to the network that introduces another loop into the VLAN can result in a broadcast storm.

Note If you have already used all available spanning-tree instances on your switch, adding another VLAN anywhere in the VTP domain creates a VLAN that is not running spanning tree on that switch. If you have the default allowed list on the trunk ports of that switch, the new VLAN is carried on all trunk ports. Depending on the topology of the network, this could create a loop in the new VLAN that will not be broken, particularly if there are several adjacent switches that have all run out of spanning-tree instances.

You can prevent this possibility by setting up allowed lists on the trunk ports of switches that have used up their allocation of spanning-tree instances. Setting up allowed lists is not necessary in many cases and can make it more labor-intensive to add another VLAN to the network.

Spanning-tree commands control the configuration of VLAN spanning-tree instances. You create a spanning-tree instance when you assign an interface to a VLAN. The spanning-tree instance is removed when the last interface is moved to another VLAN. You can configure switch and port parameters before a spanning-tree instance is created; these parameters are applied when the spanning-tree instance is created.

The switch supports PVST+, rapid PVST+, and MSTP, but only one version can be active at any time. (For example, all VLANs run PVST+, all VLANs run rapid PVST+, or all VLANs run MSTP.) For information about the different spanning-tree modes and how they interoperate, see the "Spanning-Tree Interoperability and Backward Compatibility" section on page 20-10.

For configuration information about UplinkFast and BackboneFast, see the "Information About Configuring the Optional Spanning-Tree Features" section on page 22-1.

Caution Loop guard works only on point-to-point links. We recommend that each end of the link has a directly connected device that is running STP.

How to Configure STP

Changing the Spanning-Tree Mode

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**spanning-tree mode** {**pvst** | **mst** | **rapid-pvst**}	Configures a spanning-tree mode. • **pvst**—Enables PVST+ (the default setting). • **mst**—Enables MSTP (and RSTP). For more configuration steps, see Chapter 21, "Configuring MSTP." • **rapid-pvst**—Enables rapid PVST+.
Step 3	**interface** *interface-id*	(Recommended for rapid-PVST+ mode only) Specifies an interface to configure, and enters interface configuration mode. Valid interfaces include physical ports, VLANs, and port channels.
Step 4	**spanning-tree link-type point-to-point**	(Recommended for rapid-PVST+ mode only) Specifies that the link type for this port is point-to-point. If you connect this port (local port) to a remote port through a point-to-point link and the local port becomes a designated port, the switch negotiates with the remote port and rapidly changes the local port to the forwarding state.

	Command	Purpose
Step 5	**end**	Returns to privileged EXEC mode.
Step 6	**clear spanning-tree detected-protocols**	(Recommended for rapid-PVST+ mode only) Restarts the protocol migration process on the entire switch if any port on the switch is connected to a port on a legacy IEEE 802.1D switch, This step is optional if the designated switch detects that this switch is running rapid PVST+.

Configuring the Root Switch

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**spanning-tree vlan** *vlan-id* **root primary** [**diameter** *net-diameter* [**hello-time** *seconds*]]	Configures a switch to become the root for the specified VLAN. • *vlan-id*—Specifies a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. • (Optional) **diameter** *net-diameter*—Specifies the maximum number of switches between any two end stations. • (Optional) **hello-time** *seconds*—Specifies the interval in seconds between the generation of configuration messages by the root switch.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring a Secondary Root Switch

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**spanning-tree vlan** *vlan-id* **root secondary** [**diameter** *net-diameter* [**hello-time** *seconds*]]	Configures a switch to become the secondary root for the specified VLAN. • *vlan-id*—Specifies a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. The range is 1 to 4096. • (Optional) **diameter** *net-diameter*—Specifies the maximum number of switches between any two end stations. The range is 2 to 7. • (Optional) **hello-time** *seconds*—Specifies the interval in seconds between the generation of configuration messages by the root switch. The range is 1 to 10; the default is 2. Use the same network diameter and hello-time values that you used when configuring the primary root switch. See the "Configuring the Root Switch" section on page 20-15.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring Port Priority

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies an interface to configure, and enters interface configuration mode. Valid interfaces include physical ports and port-channel logical interfaces (**port-channel** *port-channel-number*).
Step 3	**spanning-tree port-priority** *priority*	Configures the port priority for an interface.
Step 4	**spanning-tree vlan** *vlan-id* **port-priority** *priority*	Configures the port priority for a VLAN.
Step 5	**end**	Returns to privileged EXEC mode.

Configuring Path Cost

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies an interface to configure, and enters interface configuration mode. Valid interfaces include physical ports and port-channel logical interfaces (**port-channel** *port-channel-number*).

	Command	Purpose
Step 3	**spanning-tree cost** *cost*	Configures the cost for an interface.
Step 4	**spanning-tree vlan** *vlan-id* **cost** *cost*	Configures the cost for a VLAN. If a loop occurs, spanning tree uses the path cost when selecting an interface to place into the forwarding state. A lower path cost represents higher-speed transmission.
Step 5	**end**	Returns to privileged EXEC mode.

Configuring Optional STP Parameters

Before You Begin

Exercise care when configuring the priority, and hello time for STP.

For most situations, we recommend that you use the **spanning-tree vlan** *vlan-id* **root primary** and the **spanning-tree vlan** *vlan-id* **root secondary** global configuration commands to modify the switch priority.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**spanning-tree vlan** *vlan-id* **priority** *priority*	Configures the switch priority of a VLAN.
Step 3	**spanning-tree vlan** *vlan-id* **hello-time** *seconds*	Configures the hello time of a VLAN.
Step 4	**spanning-tree vlan** *vlan-id* **max-age** *seconds*	Configures the maximum-aging time of a VLAN.
Step 5	**spanning-tree vlan** *vlan-id* **forward-time** *seconds*	Configures the forward time of a VLAN.
Step 6	**spanning-tree vlan** *vlan-id* **max-age** *seconds*	Configures the maximum-aging time of a VLAN.
Step 7	**spanning-tree transmit hold-count** *value*	Configures the number of BPDUs that can be sent before pausing for 1 second. **Note** Changing this parameter to a higher value can have a significant impact on CPU utilization, especially in Rapid-PVST mode. Lowering this value can slow down convergence in certain scenarios. We recommend that you maintain the default setting.
Step 8	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining STP

Command	Purpose
show spanning-tree active	Displays spanning-tree information on active interfaces only.
show spanning-tree detail	Displays a detailed summary of interface information.
show spanning-tree interface *interface-id*	Displays spanning-tree information for the specified interface.
show spanning-tree summary	Displays a summary of interface states.

Command	Purpose
show spanning-tree vlan *vlan-id*	Displays spanning-tree VLAN entries.
copy running-config startup-config	(Optional) Saves your entries in the configuration file.

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
VLAN configuration	*Chapter 17, "Configuring VLANs"*
Multiple Spanning Tree Protocol configuration	*Chapter 21, "Configuring MSTP"*
Optional Spanning-Tree configuration	*Chapter 22, "Configuring Optional Spanning-Tree Features"*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

CHAPTER 21

Configuring MSTP

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Information About Configuring MSTP

This chapter describes how to configure the Cisco implementation of the IEEE 802.1s Multiple STP (MSTP) on the switch.

Note The multiple spanning-tree (MST) implementation is based on the IEEE 802.1s standard.

The MSTP enables multiple VLANs to be mapped to the same spanning-tree instance, reducing the number of spanning-tree instances needed to support a large number of VLANs. The MSTP provides for multiple forwarding paths for data traffic and enables load balancing. It improves the fault tolerance of the network because a failure in one instance (forwarding path) does not affect other instances (forwarding paths). The most common initial deployment of MSTP is in the backbone and distribution layers of a Layer 2 switched network. This deployment provides the highly available network required in a service-provider environment.

When the switch is in the MST mode, the Rapid Spanning Tree Protocol (RSTP), which is based on IEEE 802.1w, is automatically enabled. The RSTP provides rapid convergence of the spanning tree through explicit handshaking that eliminates the IEEE 802.1D forwarding delay and quickly transitions root ports and designated ports to the forwarding state.

Both MSTP and RSTP improve the spanning-tree operation and maintain backward compatibility with equipment that is based on the (original) IEEE 802.1D spanning tree, with existing Cisco-proprietary Multiple Instance STP (MISTP), and with existing Cisco per-VLAN spanning-tree plus (PVST+) and rapid per-VLAN spanning-tree plus (rapid PVST+).

MSTP

MSTP, which uses RSTP for rapid convergence, enables VLANs to be grouped into a spanning-tree instance, with each instance having a spanning-tree topology independent of other spanning-tree instances. This architecture provides multiple forwarding paths for data traffic, enables load balancing, and reduces the number of spanning-tree instances required to support a large number of VLANs.

Multiple Spanning-Tree Regions

For switches to participate in multiple spanning-tree (MST) instances, you must consistently configure the switches with the same MST configuration information. A collection of interconnected switches that have the same MST configuration comprises an MST region as shown in Figure 21-1 on page 21-4.

The MST configuration controls to which MST region each switch belongs. The configuration includes the name of the region, the revision number, and the MST VLAN-to-instance assignment map. You configure the switch for a region by using the **spanning-tree mst configuration** global configuration command, after which the switch enters the MST configuration mode. From this mode, you can map VLANs to an MST instance by using the **instance** MST configuration command, specify the region name by using the **name** MST configuration command, and set the revision number by using the **revision** MST configuration command.

A region can have one or multiple members with the same MST configuration. Each member must be capable of processing RSTP bridge protocol data units (BPDUs). There is no limit to the number of MST regions in a network, but each region can support up to 65 spanning-tree instances. Instances can be identified by any number in the range from 0 to 4096. You can assign a VLAN to only one spanning-tree instance at a time.

IST, CIST, and CST

Unlike PVST+ and rapid PVST+ in which all the spanning-tree instances are independent, the MSTP establishes and maintains two types of spanning trees:

- An internal spanning tree (IST), which is the spanning tree that runs in an MST region.

  Within each MST region, the MSTP maintains multiple spanning-tree instances. Instance 0 is a special instance for a region, known as the internal spanning tree (IST). All other MST instances are numbered from 1 to 4096.

  The IST is the only spanning-tree instance that sends and receives BPDUs. All of the other spanning-tree instance information is contained in M-records, which are encapsulated within MSTP BPDUs. Because the MSTP BPDU carries information for all instances, the number of BPDUs that need to be processed to support multiple spanning-tree instances is significantly reduced.

  All MST instances within the same region share the same protocol timers, but each MST instance has its own topology parameters, such as root switch ID, root path cost, and so forth. By default, all VLANs are assigned to the IST.

  An MST instance is local to the region; for example, MST instance 1 in region A is independent of MST instance 1 in region B, even if regions A and B are interconnected.

- A common and internal spanning tree (CIST), which is a collection of the ISTs in each MST region, and the common spanning tree (CST) that interconnects the MST regions and single spanning trees.

  The spanning tree computed in a region appears as a subtree in the CST that encompasses the entire switched domain. The CIST is formed by the spanning-tree algorithm running among switches that support the IEEE 802.1w, IEEE 802.1s, and IEEE 802.1D standards. The CIST inside an MST region is the same as the CST outside a region.

For more information, see the "Operations Within an MST Region" section on page 21-3 and the "Operations Between MST Regions" section on page 21-3.

Note The implementation of the IEEE 802.1s standard, changes some of the terminology associated with MST implementations. For a summary of these changes, see Table 20-1 on page 20-4.

Operations Within an MST Region

The IST connects all the MSTP switches in a region. When the IST converges, the root of the IST becomes the CIST regional root (called the *IST master* before the implementation of the IEEE 802.1s standard) as shown in Figure 21-1 on page 21-4. It is the switch within the region with the lowest switch ID and path cost to the CIST root. The CIST regional root is also the CIST root if there is only one region in the network. If the CIST root is outside the region, one of the MSTP switches at the boundary of the region is selected as the CIST regional root.

When an MSTP switch initializes, it sends BPDUs claiming itself as the root of the CIST and the CIST regional root, with both of the path costs to the CIST root and to the CIST regional root set to zero. The switch also initializes all of its MST instances and claims to be the root for all of them. If the switch receives superior MST root information (lower switch ID, lower path cost, and so forth) than currently stored for the port, it relinquishes its claim as the CIST regional root.

During initialization, a region might have many subregions, each with its own CIST regional root. As switches receive superior IST information, they leave their old subregions and join the new subregion that contains the true CIST regional root. All subregions shrink, except for the one that contains the true CIST regional root.

For correct operation, all switches in the MST region must agree on the same CIST regional root. Therefore, any two switches in the region only synchronize their port roles for an MST instance if they converge to a common CIST regional root.

Operations Between MST Regions

If there are multiple regions or legacy IEEE 802.1D switches within the network, MSTP establishes and maintains the CST, which includes all MST regions and all legacy STP switches in the network. The MST instances combine with the IST at the boundary of the region to become the CST.

The IST connects all the MSTP switches in the region and appears as a subtree in the CIST that encompasses the entire switched domain. The root of the subtree is the CIST regional root. The MST region appears as a virtual switch to adjacent STP switches and MST regions.

Figure 21-1 shows a network with three MST regions and a legacy IEEE 802.1D switch (D). The CIST regional root for region 1 (A) is also the CIST root. The CIST regional root for region 2 (B) and the CIST regional root for region 3 (C) are the roots for their respective subtrees within the CIST. The RSTP runs in all regions.

Figure 21-1 MST Regions, CIST Masters, and CST Root

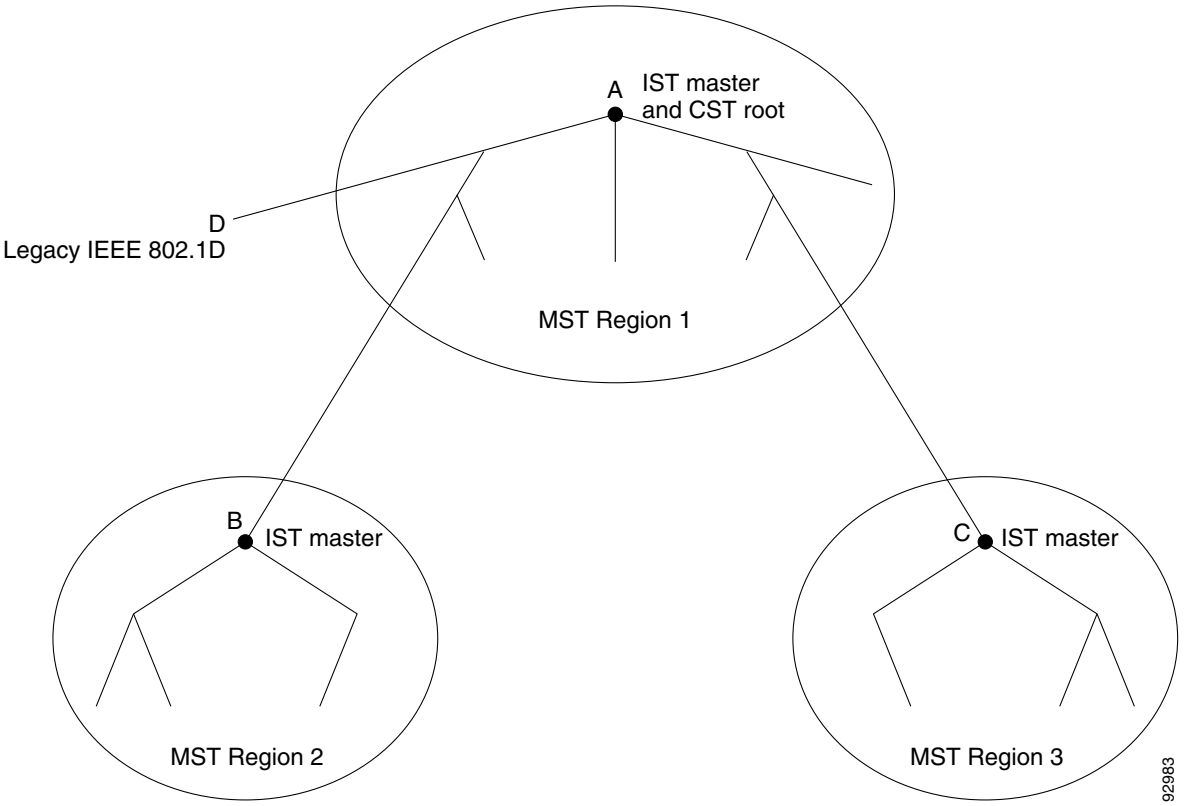

Only the CST instance sends and receives BPDUs, and MST instances add their spanning-tree information into the BPDUs to interact with neighboring switches and compute the final spanning-tree topology. Because of this, the spanning-tree parameters related to BPDU transmission (for example, hello time, forward time, max-age, and max-hops) are configured only on the CST instance but affect all MST instances. Parameters related to the spanning-tree topology (for example, switch priority, port VLAN cost, and port VLAN priority) can be configured on both the CST instance and the MST instance.

MSTP switches use Version 3 RSTP BPDUs or IEEE 802.1D STP BPDUs to communicate with legacy IEEE 802.1D switches. MSTP switches use MSTP BPDUs to communicate with MSTP switches.

IEEE 802.1s Terminology

Some MST naming conventions used in Cisco's prestandard implementation have been changed to identify some *internal* or *regional* parameters. These parameters are significant only within an MST region, as opposed to external parameters that are relevant to the whole network. Because the CIST is the only spanning-tree instance that spans the whole network, only the CIST parameters require the external rather than the internal or regional qualifiers.

- The CIST root is the root switch for the unique instance that spans the whole network, the CIST.
- The CIST external root path cost is the cost to the CIST root. This cost is left unchanged within an MST region. Remember that an MST region looks like a single switch for the CIST. The CIST external root path cost is the root path cost calculated between these virtual switches and switches that do not belong to any region.

- The CIST regional root was called the IST master in the prestandard implementation. If the CIST root is in the region, the CIST regional root is the CIST root. Otherwise, the CIST regional root is the closest switch to the CIST root in the region. The CIST regional root acts as a root switch for the IST.
- The CIST internal root path cost is the cost to the CIST regional root in a region. This cost is only relevant to the IST, instance 0.

Table 21-1 on page 21-5 compares the IEEE standard and the Cisco prestandard terminology.

Table 21-1 Prestandard and Standard Terminology

IEEE Standard	Cisco Prestandard	Cisco Standard
CIST regional root	IST master	CIST regional root
CIST internal root path cost	IST master path cost	CIST internal path cost
CIST external root path cost	Root path cost	Root path cost
MSTI regional root	Instance root	Instance root
MSTI internal root path cost	Root path cost	Root path cost

Hop Count

The IST and MST instances do not use the message-age and maximum-age information in the configuration BPDU to compute the spanning-tree topology. Instead, they use the path cost to the root and a hop-count mechanism similar to the IP time-to-live (TTL) mechanism.

By using the **spanning-tree mst max-hops** global configuration command, you can configure the maximum hops inside the region and apply it to the IST and all MST instances in that region. The hop count achieves the same result as the message-age information (triggers a reconfiguration). The root switch of the instance always sends a BPDU (or M-record) with a cost of 0 and the hop count set to the maximum value. When a switch receives this BPDU, it decrements the received remaining hop count by one and propagates this value as the remaining hop count in the BPDUs it generates. When the count reaches zero, the switch discards the BPDU and ages the information held for the port.

The message-age and maximum-age information in the RSTP portion of the BPDU remain the same throughout the region, and the same values are propagated by the region designated ports at the boundary.

Boundary Ports

In the Cisco prestandard implementation, a boundary port connects an MST region to a single spanning-tree region running RSTP, to a single spanning-tree region running PVST+ or rapid PVST+, or to another MST region with a different MST configuration. A boundary port also connects to a LAN, the designated switch of which is either a single spanning-tree switch or a switch with a different MST configuration.

There is no definition of a boundary port in the IEEE 802.1s standard. The IEEE 802.1Q-2002 standard identifies two kinds of messages that a port can receive: internal (coming from the same region) and external. When a message is external, it is received only by the CIST. If the CIST role is root or alternate, or if the external BPDU is a topology change, it could have an impact on the MST instances. When a message is internal, the CIST part is received by the CIST, and each MST instance receives its respective M-record. The Cisco prestandard implementation treats a port that receives an external message as a boundary port. This means a port cannot receive a mix of internal and external messages.

An MST region includes both switches and LANs. A segment belongs to the region of its designated port. Therefore, a port in a different region than the designated port for a segment is a boundary port. This definition allows two ports internal to a region to share a segment with a port belonging to a different region, creating the possibility of receiving both internal and external messages on a port.

The primary change from the Cisco prestandard implementation is that a designated port is not defined as boundary, unless it is running in an STP-compatible mode.

Note If there is a legacy STP switch on the segment, messages are always considered external.

The other change from the prestandard implementation is that the CIST regional root switch ID field is now inserted where an RSTP or legacy IEEE 802.1Q switch has the sender switch ID. The whole region performs like a single virtual switch by sending a consistent sender switch ID to neighboring switches. In this example, switch C would receive a BPDU with the same consistent sender switch ID of root, whether or not A or B is designated for the segment.

IEEE 802.1s Implementation

The Cisco implementation of the IEEE MST standard includes features required to meet the standard, as well as some of the desirable prestandard functionality that is not yet incorporated into the published standard.

Port Role Naming Change

The boundary role is no longer in the final MST standard, but this boundary concept is maintained in Cisco's implementation. However, an MST instance port at a boundary of the region might not follow the state of the corresponding CIST port. Two cases exist now:

- The boundary port is the root port of the CIST regional root—When the CIST instance port is proposed and is in sync, it can send back an agreement and move to the forwarding state only after all the corresponding MSTI ports are in sync (and forwarding). The MSTI ports now have a special *master* role.
- The boundary port is not the root port of the CIST regional root—The MSTI ports follow the state and role of the CIST port. The standard provides less information, and it might be difficult to understand why an MSTI port can be alternately blocking when it receives no BPDUs (MRecords). In this case, although the boundary role no longer exists, the **show** commands identify a port as boundary in the *type* column of the output.

Interoperation Between Legacy and Standard Switches

Because automatic detection of prestandard switches can fail, you can use an interface configuration command to identify prestandard ports. A region cannot be formed between a standard and a prestandard switch, but they can interoperate by using the CIST. Only the capability of load balancing over different instances is lost in that particular case. The CLI displays different flags depending on the port configuration when a port receives prestandard BPDUs. A syslog message also appears the first time a switch receives a prestandard BPDU on a port that has not been configured for prestandard BPDU transmission.

Figure 21-2 illustrates this scenario. Assume that A is a standard switch and B a prestandard switch, both configured to be in the same region. A is the root switch for the CIST, and B has a root port (BX) on segment X and an alternate port (BY) on segment Y. If segment Y flaps, and the port on BY becomes

the alternate before sending out a single prestandard BPDU, AY cannot detect that a prestandard switch is connected to Y and continues to send standard BPDUs. The port BY is fixed in a boundary, and no load balancing is possible between A and B. The same problem exists on segment X, but B might transmit topology changes.

Figure 21-2 Standard and Prestandard Switch Interoperation

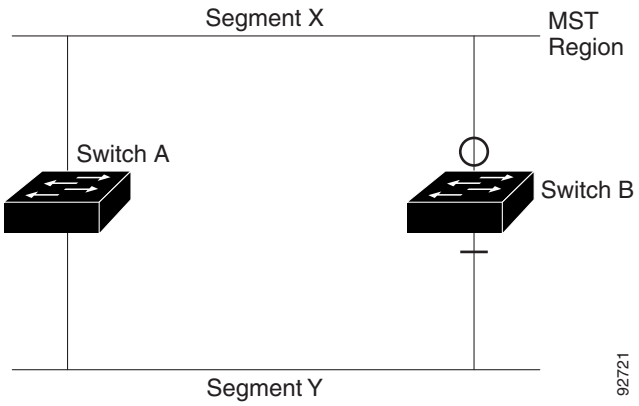

Note We recommend that you minimize the interaction between standard and prestandard MST implementations.

Detecting Unidirectional Link Failure

This feature is not yet present in the IEEE MST standard, but it is included in this Cisco IOS release. The software checks the consistency of the port role and state in the received BPDUs to detect unidirectional link failures that could cause bridging loops.

When a designated port detects a conflict, it keeps its role, but reverts to discarding state because disrupting connectivity in case of inconsistency is preferable to opening a bridging loop.

Figure 21-3 illustrates a unidirectional link failure that typically creates a bridging loop. Switch A is the root switch, and its BPDUs are lost on the link leading to switch B. RSTP and MST BPDUs include the role and state of the sending port. With this information, switch A can detect that switch B does not react to the superior BPDUs it sends and that switch B is the designated, not root switch. As a result, switch A blocks (or keeps blocking) its port, preventing the bridging loop.

Figure 21-3 Detecting Unidirectional Link Failure

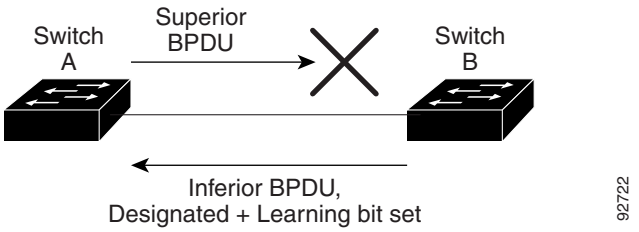

Interoperability with IEEE 802.1D STP

A switch running MSTP supports a built-in protocol migration mechanism that enables it to interoperate with legacy IEEE 802.1D switches. If this switch receives a legacy IEEE 802.1D configuration BPDU (a BPDU with the protocol version set to 0), it sends only IEEE 802.1D BPDUs on that port. An MSTP switch also can detect that a port is at the boundary of a region when it receives a legacy BPDU, an MSTP BPDU (Version 3) associated with a different region, or an RSTP BPDU (Version 2).

However, the switch does not automatically revert to the MSTP mode if it no longer receives IEEE 802.1D BPDUs because it cannot detect whether the legacy switch has been removed from the link unless the legacy switch is the designated switch. A switch might also continue to assign a boundary role to a port when the switch to which this switch is connected has joined the region. To restart the protocol migration process (force the renegotiation with neighboring switches), use the **clear spanning-tree detected-protocols** privileged EXEC command.

If all the legacy switches on the link are RSTP switches, they can process MSTP BPDUs as if they are RSTP BPDUs. Therefore, MSTP switches send either a Version 0 configuration and TCN BPDUs or Version 3 MSTP BPDUs on a boundary port. A boundary port connects to a LAN, the designated switch of which is either a single spanning-tree switch or a switch with a different MST configuration.

RSTP

The RSTP takes advantage of point-to-point wiring and provides rapid convergence of the spanning tree. Reconfiguration of the spanning tree can occur in less than 1 second (in contrast to 50 seconds with the default settings in the IEEE 802.1D spanning tree).

Port Roles and the Active Topology

The RSTP provides rapid convergence of the spanning tree by assigning port roles and by learning the active topology. The RSTP builds upon the IEEE 802.1D STP to select the switch with the highest switch priority (lowest numerical priority value) as the root switch as described in the "Spanning-Tree Topology and BPDUs" section on page 20-2. Then the RSTP assigns one of these port roles to individual ports:

- Root port—Provides the best path (lowest cost) when the switch forwards packets to the root switch.
- Designated port—Connects to the designated switch, which incurs the lowest path cost when forwarding packets from that LAN to the root switch. The port through which the designated switch is attached to the LAN is called the designated port.
- Alternate port—Offers an alternate path toward the root switch to that provided by the current root port.
- Backup port—Acts as a backup for the path provided by a designated port toward the leaves of the spanning tree. A backup port can exist only when two ports are connected in a loopback by a point-to-point link or when a switch has two or more connections to a shared LAN segment.
- Disabled port—Has no role within the operation of the spanning tree.

A port with the root or a designated port role is included in the active topology. A port with the alternate or backup port role is excluded from the active topology.

In a stable topology with consistent port roles throughout the network, the RSTP ensures that every root port and designated port immediately transition to the forwarding state while all alternate and backup ports are always in the discarding state (equivalent to blocking in IEEE 802.1D). The port state controls the operation of the forwarding and learning processes. Table 21-2 provides a comparison of IEEE 802.1D and RSTP port states.

Table 21-2 Port State Comparison

Operational Status	STP Port State (IEEE 802.1D)	RSTP Port State	Is Port Included in the Active Topology?
Enabled	Blocking	Discarding	No
Enabled	Listening	Discarding	No
Enabled	Learning	Learning	Yes
Enabled	Forwarding	Forwarding	Yes
Disabled	Disabled	Discarding	No

To be consistent with Cisco STP implementations, this guide defines the port state as *blocking* instead of *discarding*. Designated ports start in the listening state.

Rapid Convergence

The RSTP provides for rapid recovery of connectivity following the failure of a switch, a switch port, or a LAN. It provides rapid convergence for edge ports, new root ports, and ports connected through point-to-point links as follows:

- Edge ports—If you configure a port as an edge port on an RSTP switch by using the **spanning-tree portfast** interface configuration command, the edge port immediately transitions to the forwarding state. An edge port is the same as a Port Fast-enabled port, and you should enable it only on ports that connect to a single end station.
- Root ports—If the RSTP selects a new root port, it blocks the old root port and immediately transitions the new root port to the forwarding state.
- Point-to-point links—If you connect a port to another port through a point-to-point link and the local port becomes a designated port, it negotiates a rapid transition with the other port by using the proposal-agreement handshake to ensure a loop-free topology.

  As shown in Figure 21-4, Switch A is connected to Switch B through a point-to-point link, and all of the ports are in the blocking state. Assume that the priority of Switch A is a smaller numerical value than the priority of Switch B. Switch A sends a proposal message (a configuration BPDU with the proposal flag set) to Switch B, proposing itself as the designated switch.

  After receiving the proposal message, Switch B selects as its new root port the port from which the proposal message was received, forces all nonedge ports to the blocking state, and sends an agreement message (a BPDU with the agreement flag set) through its new root port.

  After receiving Switch B's agreement message, Switch A also immediately transitions its designated port to the forwarding state. No loops in the network are formed because Switch B blocked all of its nonedge ports and because there is a point-to-point link between Switches A and B.

When Switch C is connected to Switch B, a similar set of handshaking messages are exchanged. Switch C selects the port connected to Switch B as its root port, and both ends immediately transition to the forwarding state. With each iteration of this handshaking process, one more switch joins the active topology. As the network converges, this proposal-agreement handshaking progresses from the root toward the leaves of the spanning tree.

The switch learns the link type from the port duplex mode: a full-duplex port is considered to have a point-to-point connection; a half-duplex port is considered to have a shared connection. You can override the default setting that is controlled by the duplex setting by using the **spanning-tree link-type** interface configuration command.

Figure 21-4 Proposal and Agreement Handshaking for Rapid Convergence

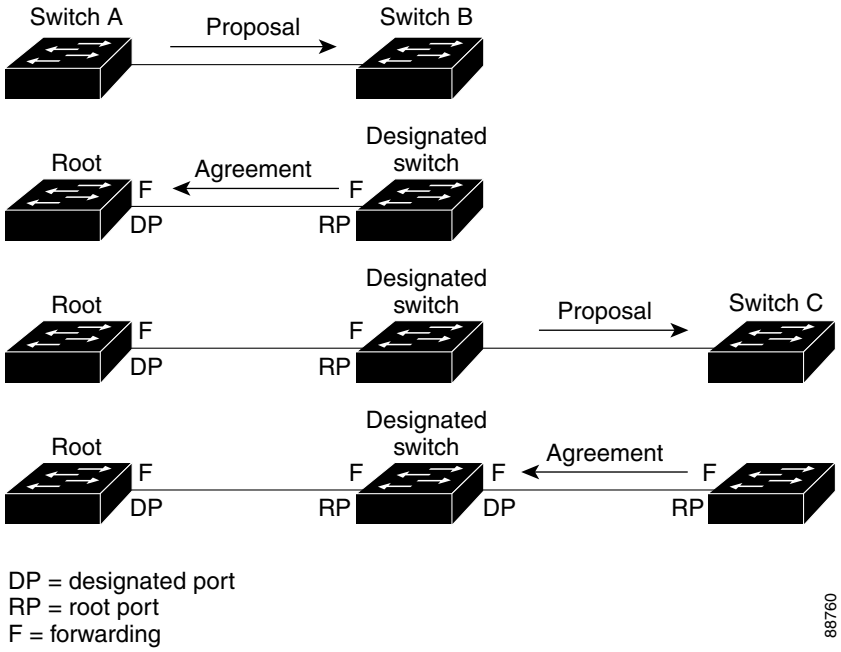

Synchronization of Port Roles

When the switch receives a proposal message on one of its ports and that port is selected as the new root port, the RSTP forces all other ports to synchronize with the new root information.

The switch is synchronized with superior root information received on the root port if all other ports are synchronized. An individual port on the switch is synchronized if

- That port is in the blocking state.
- It is an edge port (a port configured to be at the edge of the network).

If a designated port is in the forwarding state and is not configured as an edge port, it transitions to the blocking state when the RSTP forces it to synchronize with new root information. In general, when the RSTP forces a port to synchronize with root information and the port does not satisfy any of the above conditions, its port state is set to blocking.

After ensuring that all of the ports are synchronized, the switch sends an agreement message to the designated switch corresponding to its root port. When the switches connected by a point-to-point link are in agreement about their port roles, the RSTP immediately transitions the port states to forwarding. The sequence of events is shown in Figure 21-5.

Figure 21-5 Sequence of Events During Rapid Convergence

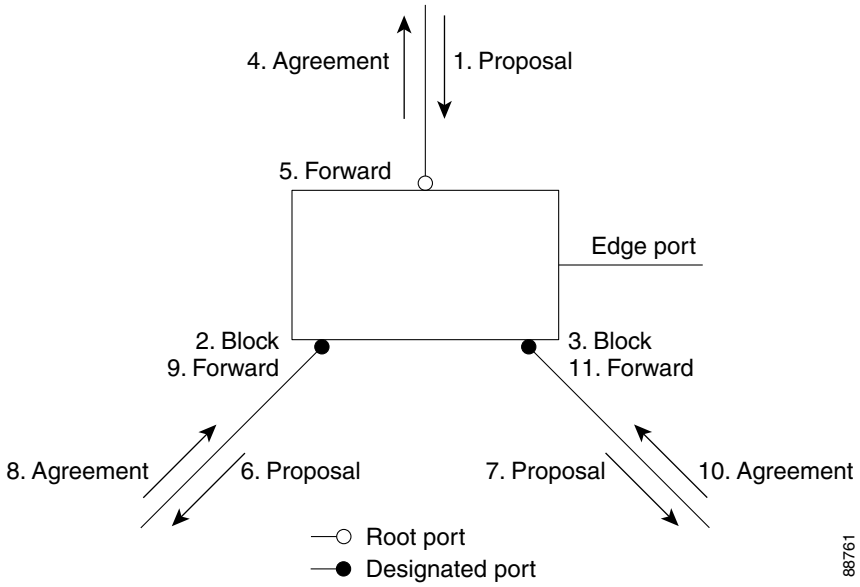

Bridge Protocol Data Unit Format and Processing

The RSTP BPDU format is the same as the IEEE 802.1D BPDU format except that the protocol version is set to 2. A new 1-byte Version 1 Length field is set to zero, which means that no version 1 protocol information is present. Table 21-3 shows the RSTP flag fields.

Table 21-3 RSTP BPDU Flags

Bit	Function
0	Topology change (TC)
1	Proposal
2–3:	Port role:
00	Unknown
01	Alternate port
10	Root port
11	Designated port
4	Learning
5	Forwarding
6	Agreement
7	Topology change acknowledgement (TCA)

The sending switch sets the proposal flag in the RSTP BPDU to propose itself as the designated switch on that LAN. The port role in the proposal message is always set to the designated port.

The sending switch sets the agreement flag in the RSTP BPDU to accept the previous proposal. The port role in the agreement message is always set to the root port.

The RSTP does not have a separate topology change notification (TCN) BPDU. It uses the topology change (TC) flag to show the topology changes. However, for interoperability with IEEE 802.1D switches, the RSTP switch processes and generates TCN BPDUs.

The learning and forwarding flags are set according to the state of the sending port.

Processing Superior BPDU Information

If a port receives superior root information (lower switch ID, lower path cost, and so forth) than currently stored for the port, the RSTP triggers a reconfiguration. If the port is proposed and is selected as the new root port, RSTP forces all the other ports to synchronize.

If the BPDU received is an RSTP BPDU with the proposal flag set, the switch sends an agreement message after all of the other ports are synchronized. If the BPDU is an IEEE 802.1D BPDU, the switch does not set the proposal flag and starts the forward-delay timer for the port. The new root port requires twice the forward-delay time to transition to the forwarding state.

If the superior information received on the port causes the port to become a backup or alternate port, RSTP sets the port to the blocking state but does not send the agreement message. The designated port continues sending BPDUs with the proposal flag set until the forward-delay timer expires, at which time the port transitions to the forwarding state.

Processing Inferior BPDU Information

If a designated port receives an inferior BPDU (higher switch ID, higher path cost, and so forth than currently stored for the port) with a designated port role, it immediately replies with its own information.

Topology Changes

This section describes the differences between the RSTP and the IEEE 802.1D in handling spanning-tree topology changes.

- Detection—Unlike IEEE 802.1D in which *any* transition between the blocking and the forwarding state causes a topology change, *only* transitions from the blocking to the forwarding state cause a topology change with RSTP (only an increase in connectivity is considered a topology change). State changes on an edge port do not cause a topology change. When an RSTP switch detects a topology change, it deletes the learned information on all of its nonedge ports except on those from which it received the TC notification.
- Notification—Unlike IEEE 802.1D, which uses TCN BPDUs, the RSTP does not use them. However, for IEEE 802.1D interoperability, an RSTP switch processes and generates TCN BPDUs.
- Acknowledgement—When an RSTP switch receives a TCN message on a designated port from an IEEE 802.1D switch, it replies with an IEEE 802.1D configuration BPDU with the TCA bit set. However, if the TC-while timer (the same as the topology-change timer in IEEE 802.1D) is active on a root port connected to an IEEE 802.1D switch and a configuration BPDU with the TCA bit set is received, the TC-while timer is reset.

  This behavior is only required to support IEEE 802.1D switches. The RSTP BPDUs never have the TCA bit set.

- Propagation—When an RSTP switch receives a TC message from another switch through a designated or root port, it propagates the change to all of its nonedge, designated ports and to the root port (excluding the port on which it is received). The switch starts the TC-while timer for all such ports and flushes the information learned on them.
- Protocol migration—For backward compatibility with IEEE 802.1D switches, RSTP selectively sends IEEE 802.1D configuration BPDUs and TCN BPDUs on a per-port basis.

  When a port is initialized, the migrate-delay timer is started (specifies the minimum time during which RSTP BPDUs are sent), and RSTP BPDUs are sent. While this timer is active, the switch processes all BPDUs received on that port and ignores the protocol type.

  If the switch receives an IEEE 802.1D BPDU after the port migration-delay timer has expired, it assumes that it is connected to an IEEE 802.1D switch and starts using only IEEE 802.1D BPDUs. However, if the RSTP switch is using IEEE 802.1D BPDUs on a port and receives an RSTP BPDU after the timer has expired, it restarts the timer and starts using RSTP BPDUs on that port.

Default MSTP Settings

Table 21-4 Default MSTP Settings

Feature	Default Setting
Spanning-tree mode	PVST+ (Rapid PVST+ and MSTP are disabled)
Switch priority (configurable on a per-CIST port basis)	32768
Spanning-tree port priority (configurable on a per-CIST port basis)	128
Spanning-tree port cost (configurable on a per-CIST port basis)	1000 Mbps: 4 100 Mbps: 19 10 Mbps: 100
Hello time	2 seconds
Forward-delay time	15 seconds
Maximum-aging time	20 seconds
Maximum hop count	20 hops

MSTP Configuration Guidelines

These are the configuration guidelines for MSTP:

- When you enable MST by using the **spanning-tree mode mst** global configuration command, RSTP is automatically enabled.
- For two or more switches to be in the same MST region, they must have the same VLAN-to-instance map, the same configuration revision number, and the same name.
- The switch supports up to 65 MST instances. The number of VLANs that can be mapped to a particular MST instance is unlimited.
- PVST+, rapid PVST+, and MSTP are supported, but only one version can be active at any time. (For example, all VLANs run PVST+, all VLANs run rapid PVST+, or all VLANs run MSTP.) For more information, see the "Spanning-Tree Interoperability and Backward Compatibility" section on page 20-10.

- VTP propagation of the MST configuration is not supported. However, you can manually configure the MST configuration (region name, revision number, and VLAN-to-instance mapping) on each switch within the MST region by using the command-line interface (CLI) or through the SNMP support.
- For load balancing across redundant paths in the network to work, all VLAN-to-instance mapping assignments must match; otherwise, all traffic flows on a single link.
- All MST boundary ports must be forwarding for load balancing between a PVST+ and an MST cloud or between a rapid-PVST+ and an MST cloud. For this to occur, the IST master of the MST cloud should also be the root of the CST. If the MST cloud consists of multiple MST regions, one of the MST regions must contain the CST root, and all of the other MST regions must have a better path to the root contained within the MST cloud than a path through the PVST+ or rapid-PVST+ cloud. You might have to manually configure the switches in the clouds.
- Partitioning the network into a large number of regions is not recommended. However, if this situation is unavoidable, we recommend that you partition the switched LAN into smaller LANs interconnected by routers or non-Layer 2 devices.
- For configuration information about UplinkFast and BackboneFast, see the "Information About Configuring the Optional Spanning-Tree Features" section on page 22-1.

Root Switch

The switch maintains a spanning-tree instance for the group of VLANs mapped to it. A switch ID, consisting of the switch priority and the switch MAC address, is associated with each instance. For a group of VLANs, the switch with the lowest switch ID becomes the root switch.

To configure a switch to become the root, use the **spanning-tree mst** *instance-id* **root** global configuration command to modify the switch priority from the default value (32768) to a significantly lower value so that the switch becomes the root switch for the specified spanning-tree instance. When you enter this command, the switch checks the switch priorities of the root switches. Because of the extended system ID support, the switch sets its own priority for the specified instance to 24576 if this value will cause this switch to become the root for the specified spanning-tree instance.

If any root switch for the specified instance has a switch priority lower than 24576, the switch sets its own priority to 4096 less than the lowest switch priority. (4096 is the value of the least-significant bit of a 4-bit switch priority value as shown in Table 20-1 on page 20-4.)

If your network consists of switches that both do and do not support the extended system ID, it is unlikely that the switch with the extended system ID support will become the root switch. The extended system ID increases the switch priority value every time the VLAN number is greater than the priority of the connected switches running older software.

The root switch for each spanning-tree instance should be a backbone or distribution switch. Do not configure an access switch as the spanning-tree primary root.

Use the **diameter** keyword, which is available only for MST instance 0, to specify the Layer 2 network diameter (that is, the maximum number of switch hops between any two end stations in the Layer 2 network). When you specify the network diameter, the switch automatically sets an optimal hello time, forward-delay time, and maximum-age time for a network of that diameter, which can significantly reduce the convergence time. You can use the **hello** keyword to override the automatically calculated hello time.

Secondary Root Switch

When you configure a switch with the extended system ID support as the secondary root, the switch priority is modified from the default value (32768) to 28672. The switch is then likely to become the root switch for the specified instance if the primary root switch fails. This is assuming that the other network switches use the default switch priority of 32768 and therefore are unlikely to become the root switch.

You can execute this command on more than one switch to configure multiple backup root switches. Use the same network diameter and hello-time values that you used when you configured the primary root switch with the **spanning-tree mst** *instance-id* **root primary** global configuration command.

Port Priority

If a loop occurs, the MSTP uses the port priority when selecting an interface to put into the forwarding state. You can assign higher priority values (lower numerical values) to interfaces that you want selected first and lower priority values (higher numerical values) that you want selected last. If all interfaces have the same priority value, the MSTP puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.

Path Cost

The MSTP path cost default value is derived from the media speed of an interface. If a loop occurs, the MSTP uses cost when selecting an interface to put in the forwarding state. You can assign lower cost values to interfaces that you want selected first and higher cost values that you want selected last. If all interfaces have the same cost value, the MSTP puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.

Link Type to Ensure Rapid Transitions

If you connect a port to another port through a point-to-point link and the local port becomes a designated port, the RSTP negotiates a rapid transition with the other port by using the proposal-agreement handshake to ensure a loop-free topology as described in the "Rapid Convergence" section on page 21-9.

By default, the link type is controlled from the duplex mode of the interface: a full-duplex port is considered to have a point-to-point connection; a half-duplex port is considered to have a shared connection. If you have a half-duplex link physically connected point-to-point to a single port on a remote switch running MSTP, you can override the default setting of the link type and enable rapid transitions to the forwarding state.

Neighbor Type

A topology could contain both prestandard and IEEE 802.1s standard compliant devices. By default, ports can automatically detect prestandard devices, but they can still receive both standard and prestandard BPDUs. When there is a mismatch between a device and its neighbor, only the CIST runs on the interface.

You can choose to set a port to send only prestandard BPDUs. The prestandard flag appears in all the show commands, even if the port is in STP compatibility mode.

Restarting the Protocol Migration Process

A switch running MSTP supports a built-in protocol migration mechanism that enables it to interoperate with legacy IEEE 802.1D switches. If this switch receives a legacy IEEE 802.1D configuration BPDU (a BPDU with the protocol version set to 0), it sends only IEEE 802.1D BPDUs on that port. An MSTP switch also can detect that a port is at the boundary of a region when it receives a legacy BPDU, an MST BPDU (Version 3) associated with a different region, or an RST BPDU (Version 2).

However, the switch does not automatically revert to the MSTP mode if it no longer receives IEEE 802.1D BPDUs because it cannot detect whether the legacy switch has been removed from the link unless the legacy switch is the designated switch. A switch also might continue to assign a boundary role to a port when the switch to which it is connected has joined the region.

How to Configure MSTP

Specifying the MST Region Configuration and Enabling MSTP

This task is required.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**spanning-tree mst configuration**	Enters MST configuration mode.
Step 3	**instance** *instance-id* **vlan** *vlan-range*	Maps VLANs to an MST instance. • *instance-id*—range is 0 to 4096. • **vlan** *vlan-range*—range is 1 to 4096. When you map VLANs to an MST instance, the mapping is incremental, and the VLANs specified in the command are added to or removed from the VLANs that were previously mapped. To specify a VLAN range, use a hyphen; for example, **instance 1 vlan 1-63** maps VLANs 1 through 63 to MST instance 1. To specify a VLAN series, use a comma; for example, **instance 1 vlan 10, 20, 30** maps VLANs 10, 20, and 30 to MST instance 1.
Step 4	**name** *name*	Specifies the configuration name. The *name* string has a maximum length of 32 characters and is case sensitive.
Step 5	**revision** *version*	Specifies the configuration revision number. The range is 0 to 65535.
Step 6	**show pending**	Verifies your configuration by displaying the pending configuration.
Step 7	**exit**	Applies all changes, and returns to global configuration mode.

	Command	Purpose
Step 8	**spanning-tree mode mst**	Enables MSTP. RSTP is also enabled. **Caution** Changing spanning-tree modes can disrupt traffic because all spanning-tree instances are stopped for the previous mode and restarted in the new mode. You cannot run both MSTP and PVST+ or both MSTP and rapid PVST+ at the same time.
Step 9	**end**	Returns to privileged EXEC mode.

Configuring the Root Switch

Before You Begin

After configuring the switch as the root switch, we recommend that you avoid manually configuring the hello time, forward-delay time, and maximum-age time through the **spanning-tree mst hello-time**, **spanning-tree mst forward-time**, and the **spanning-tree mst max-age** global configuration commands.

This task is optional.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**spanning-tree mst** *instance-id* **root primary** [**diameter** *net-diameter* [**hello-time** *seconds*]]	Configures a switch as the root switch. • *instance-id*—Specifies a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 4096. • (Optional) **diameter** *net-diameter*—Specifies the maximum number of switches between any two end stations. The range is 2 to 7. This keyword is available only for MST instance 0. • (Optional) **hello-time** *seconds*—Specifies the interval in seconds between the generation of configuration messages by the root switch. The range is 1 to 10 seconds; the default is 2 seconds.

	Command	Purpose
Step 3	**spanning-tree mst** *instance-id* **root secondary** [**diameter** *net-diameter* [**hello-time** *seconds*]]	Configures a switch as the secondary root switch. • *instance-id*—Specifies a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 4096. • (Optional) **diameter** *net-diameter*—Specifies the maximum number of switches between any two end stations. The range is 2 to 7. This keyword is available only for MST instance 0. • (Optional) **hello-time** *seconds*—Specifies the interval in seconds between the generation of configuration messages by the root switch. The range is 1 to 10 seconds; the default is 2 seconds. Use the same network diameter and hello-time values that you used when configuring the primary root switch.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring the Optional MSTP Parameters

Before You Begin

Exercise care when configuring the switch priority. For most situations, we recommend that you use the **spanning-tree mst** *instance-id* **root primary** and the **spanning-tree mst** *instance-id* **root secondary** global configuration commands to modify the switch priority.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**spanning-tree mst** *instance-id* **priority** *priority*	Configures the switch priority. • *instance-id*—Specifies a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 4096. • *priority*—The range is 0 to 61440 in increments of 4096; the default is 32768. The lower the number, the more likely the switch will be chosen as the root switch. Priority values are 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440. All other values are rejected.
Step 3	**spanning-tree mst hello-time** *seconds*	Configures the hello time for all MST instances. The hello time is the interval between the generation of configuration messages by the root switch. These messages mean that the switch is alive. *seconds*—The range is 1 to 10; the default is 2.

	Command	Purpose
Step 4	**spanning-tree mst forward-time** *seconds*	Configures the forward time for all MST instances. The forward delay is the number of seconds a port waits before changing from its spanning-tree learning and listening states to the forwarding state. *seconds*—The range is 4 to 30; the default is 15.
Step 5	**spanning-tree mst max-age** *seconds*	Configures the maximum-aging time for all MST instances. The maximum-aging time is the number of seconds a switch waits without receiving spanning-tree configuration messages before attempting a reconfiguration. *seconds*—The range is 6 to 40; the default is 20.
Step 6	**spanning-tree mst max-hops** *hop-count*	Specifies the number of hops in a region before the BPDU is discarded, and the information held for a port is aged. *hop-count*—The range is 1 to 255; the default is 20.
Step 7	**interface** *interface-id*	Specifies an interface to configure, and enters interface configuration mode. Valid interfaces include physical ports and port-channel logical interfaces.
Step 8	**spanning-tree mst** *instance-id* **port-priority** *priority*	Configures the port priority. • *instance-id*—Specifies a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 4096. • *priority*—The range is 0 to 240 in increments of 16. The default is 128. The lower the number, the higher the priority. The priority values are 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 192, 208, 224, and 240. All other values are rejected.
Step 9	**spanning-tree mst** *instance-id* **cost** *cost*	Configures the cost. If a loop occurs, the MSTP uses the path cost when selecting an interface to place into the forwarding state. A lower path cost represents higher-speed transmission. • *instance-id*—Specifies a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 4096. • *cost*—The range is 1 to 200000000; the default value is derived from the media speed of the interface.
Step 10	**spanning-tree link-type point-to-point**	Specifies that the link type of a port is point-to-point.
Step 11	**spanning-tree mst pre-standard**	Specifies that the port can send only prestandard BPDUs.
Step 12	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining MSTP

Command	Purpose
show spanning-tree mst configuration	Displays the MST region configuration.
show spanning-tree mst configuration digest	Displays the MD5 digest included in the current MSTCI.
show spanning-tree mst *instance-id*	Displays MST information for the specified instance.
show spanning-tree mst interface *interface-id*	Displays MST information for the specified interface.
clear spanning-tree detected-protocols	Restarts the protocol migration process (forces the renegotiation with neighboring switches) on the switch,
clear spanning-tree detected-protocols interface *interface-id*	Restarts the protocol migration process on a specific interface.
show running-config	Verifies your entries.
copy running-config startup-config	Saves your entries in the configuration file.

Configuration Examples for Configuring MSTP

Configuring the MST Region: Example

This example shows how to enter MST configuration mode, map VLANs 10 to 20 to MST instance 1, name the region *region1*, set the configuration revision to 1, display the pending configuration, apply the changes, and return to global configuration mode:

```
Switch(config)# spanning-tree mst configuration
Switch(config-mst)# instance 1 vlan 10-20
Switch(config-mst)# name region1
Switch(config-mst)# revision 1
Switch(config-mst)# show pending
Pending MST configuration
Name      [region1]
Revision  1
Instance  Vlans Mapped
--------  --------------------
0         1-9,21-4096
1         10-20
-------------------------------

Switch(config-mst)# exit
Switch(config)#
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
PVST+ and rapid PVST+ configuration	*Chapter 17, "Configuring VLANs"*
Optional Spanning-Tree configuration	*Chapter 22, "Configuring Optional Spanning-Tree Features"*
Supported number of spanning-tree instances	*Chapter 20, "Supported Spanning-Tree Instances"*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

CHAPTER 22

Configuring Optional Spanning-Tree Features

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for the Optional Spanning-Tree Features

You can configure all of these features when your switch is running the per-VLAN spanning-tree plus (PVST+). You can configure only the noted features when your switch is running the Multiple Spanning Tree Protocol (MSTP) or the rapid per-VLAN spanning-tree plus (rapid-PVST+) protocol.

Restrictions for the Optional Spanning-Tree Features

You can configure the UplinkFast or the BackboneFast feature for rapid PVST+ or for the MSTP, but the feature remains disabled (inactive) until you change the spanning-tree mode to PVST+.

Information About Configuring the Optional Spanning-Tree Features

PortFast

PortFast immediately brings an interface configured as an access or trunk port to the forwarding state from a blocking state, bypassing the listening and learning states. You can use PortFast on interfaces connected to a single workstation or server, as shown in Figure 22-1, to allow those devices to immediately connect to the network, rather than waiting for the spanning tree to converge.

Interfaces connected to a single workstation or server should not receive bridge protocol data units (BPDUs). An interface with PortFast enabled goes through the normal cycle of spanning-tree status changes when the switch is restarted.

Note Because the purpose of PortFast is to minimize the time interfaces must wait for spanning-tree to converge, it is effective only when used on interfaces connected to end stations. If you enable PortFast on an interface connecting to another switch, you risk creating a spanning-tree loop.

You can enable this feature by using the **spanning-tree portfast** interface configuration or the **spanning-tree portfast default** global configuration command.

Figure 22-1 PortFast-Enabled Interfaces

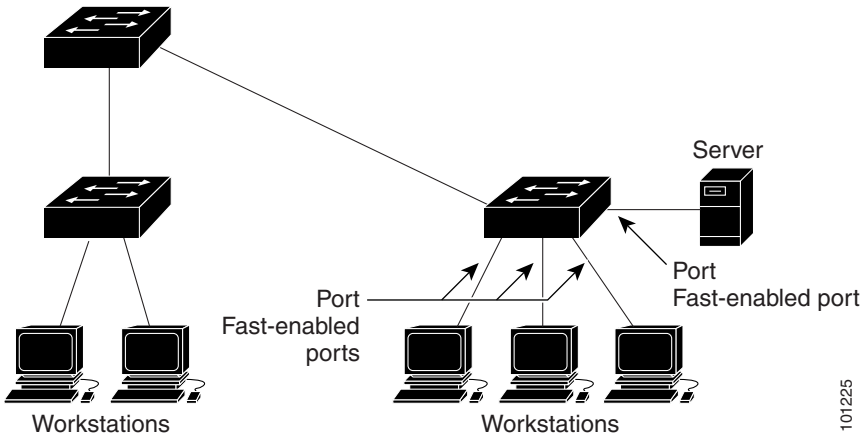

BPDU Guard

The BPDU guard feature can be globally enabled on the switch or can be enabled per port, but the feature operates with some differences.

At the global level, you enable BPDU guard on PortFast-enabled ports by using the **spanning-tree portfast bpduguard default** global configuration command. Spanning tree shuts down ports that are in a PortFast-operational state if any BPDU is received on them. In a valid configuration, PortFast-enabled ports do not receive BPDUs. Receiving a BPDU on a PortFast-enabled port means an invalid configuration, such as the connection of an unauthorized device, and the BPDU guard feature puts the port in the error-disabled state. When this happens, the switch shuts down the entire port on which the violation occurred.

To prevent the port from shutting down, you can use the **errdisable detect cause bpduguard shutdown vlan** global configuration command to shut down just the offending VLAN on the port where the violation occurred.

At the interface level, you enable BPDU guard on any port by using the **spanning-tree bpduguard enable** interface configuration command without also enabling the PortFast feature. When the port receives a BPDU, it is put in the error-disabled state.

The BPDU guard feature provides a secure response to invalid configurations because you must manually put the interface back in service. Use the BPDU guard feature in a service-provider network to prevent an access port from participating in the spanning tree.

BPDU Filtering

The BPDU filtering feature can be globally enabled on the switch or can be enabled per interface, but the feature operates with some differences.

At the global level, you can enable BPDU filtering on PortFast-enabled interfaces by using the **spanning-tree portfast bpdufilter default** global configuration command. This command prevents interfaces that are in a PortFast-operational state from sending or receiving BPDUs. The interfaces still send a few BPDUs at link-up before the switch begins to filter outbound BPDUs. You should globally enable BPDU filtering on a switch so that hosts connected to these interfaces do not receive BPDUs. If a BPDU is received on a PortFast-enabled interface, the interface loses its PortFast-operational status, and BPDU filtering is disabled.

At the interface level, you can enable BPDU filtering on any interface by using the **spanning-tree bpdufilter enable** interface configuration command without also enabling the PortFast feature. This command prevents the interface from sending or receiving BPDUs.

Caution Enabling BPDU filtering on an interface is the same as disabling spanning tree on it and can result in spanning-tree loops.

You can enable the BPDU filtering feature for the entire switch or for an interface.

UplinkFast

Switches in hierarchical networks can be grouped into backbone switches, distribution switches, and access switches. Figure 22-2 shows a complex network where distribution switches and access switches each have at least one redundant link that spanning tree blocks to prevent loops.

Figure 22-2 Switches in a Hierarchical Network

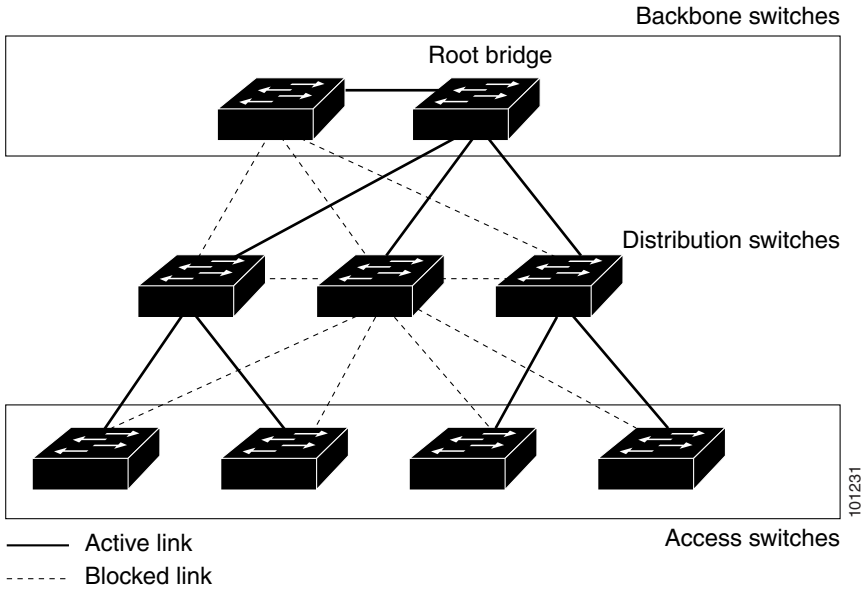

If a switch loses connectivity, it begins using the alternate paths as soon as the spanning tree selects a new root port. By enabling UplinkFast with the **spanning-tree uplinkfast** global configuration command, you can accelerate the choice of a new root port when a link or switch fails or when the spanning tree reconfigures itself. The root port transitions to the forwarding state immediately without going through the listening and learning states, as it would with the normal spanning-tree procedures.

When the spanning tree reconfigures the new root port, other interfaces flood the network with multicast packets, one for each address that was learned on the interface. You can limit these bursts of multicast traffic by reducing the max-update-rate parameter (the default for this parameter is 150 packets per second). However, if you enter zero, station-learning frames are not generated, so the spanning-tree topology converges more slowly after a loss of connectivity.

Note UplinkFast is most useful in wiring-closet switches at the access or edge of the network. It is not appropriate for backbone devices. This feature might not be useful for other types of applications.

UplinkFast provides fast convergence after a direct link failure and achieves load balancing between redundant Layer 2 links using uplink groups. An uplink group is a set of Layer 2 interfaces (per VLAN), only one of which is forwarding at any given time. Specifically, an uplink group consists of the root port (which is forwarding) and a set of blocked ports, except for self-looping ports. The uplink group provides an alternate path in case the currently forwarding link fails.

Figure 22-3 shows an example topology with no link failures. Switch A, the root switch, is connected directly to Switch B over link L1 and to Switch C over link L2. The Layer 2 interface on Switch C that is connected directly to Switch B is in a blocking state.

Figure 22-3 UplinkFast Example Before Direct Link Failure

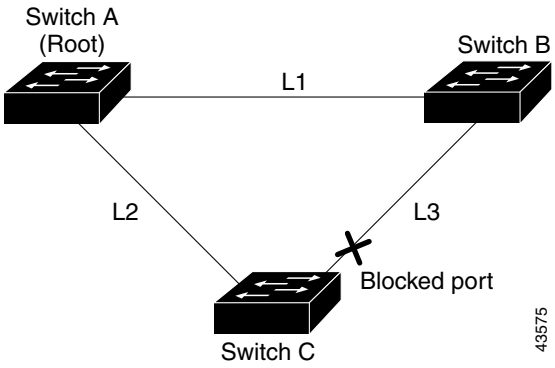

If Switch C detects a link failure on the currently active link L2 on the root port (a *direct* link failure), UplinkFast unblocks the blocked interface on Switch C and transitions it to the forwarding state without going through the listening and learning states, as shown in Figure 22-4. This change takes approximately 1 to 5 seconds.

Figure 22-4 UplinkFast Example After Direct Link Failure

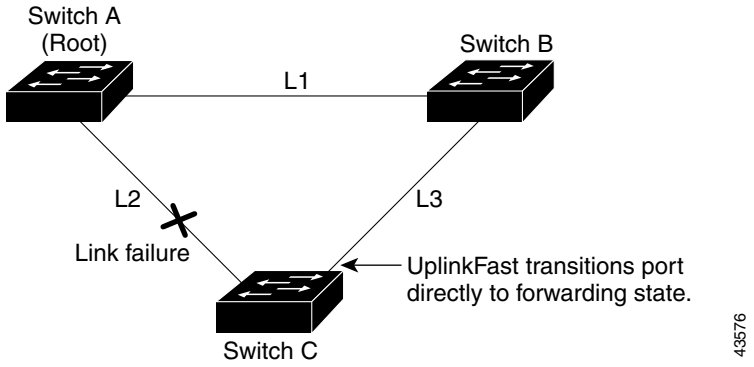

BackboneFast

BackboneFast detects indirect failures in the core of the backbone. BackboneFast is a complementary technology to the UplinkFast feature, which responds to failures on links directly connected to access switches. BackboneFast optimizes the maximum-age timer, which controls the amount of time the switch stores protocol information received on an interface. When a switch receives an inferior BPDU from the designated port of another switch, the BPDU is a signal that the other switch might have lost its path to the root, and BackboneFast tries to find an alternate path to the root.

BackboneFast, which is enabled by using the **spanning-tree backbonefast** global configuration command, starts when a root port or blocked interface on a switch receives inferior BPDUs from its designated switch. An inferior BPDU identifies a switch that declares itself as both the root bridge and the designated switch. When a switch receives an inferior BPDU, it means that a link to which the switch is not directly connected (an *indirect* link) has failed (that is, the designated switch has lost its connection to the root switch). Under spanning-tree rules, the switch ignores inferior BPDUs for the configured maximum aging time specified by the **spanning-tree vlan** *vlan-id max-age* global configuration command.

The switch tries to find if it has an alternate path to the root switch. If the inferior BPDU arrives on a blocked interface, the root port and other blocked interfaces on the switch become alternate paths to the root switch. (Self-looped ports are not considered alternate paths to the root switch.) If the inferior BPDU arrives on the root port, all blocked interfaces become alternate paths to the root switch. If the inferior BPDU arrives on the root port and there are no blocked interfaces, the switch assumes that it has lost connectivity to the root switch, causes the maximum aging time on the root port to expire, and becomes the root switch according to normal spanning-tree rules.

If the switch has alternate paths to the root switch, it uses these alternate paths to send a root link query (RLQ) request. The switch sends the RLQ request on all alternate paths and waits for an RLQ reply from other switches in the network.

If the switch discovers that it still has an alternate path to the root, it expires the maximum aging time on the interface that received the inferior BPDU. If all the alternate paths to the root switch indicate that the switch has lost connectivity to the root switch, the switch expires the maximum aging time on the interface that received the RLQ reply. If one or more alternate paths can still connect to the root switch, the switch makes all interfaces on which it received an inferior BPDU its designated ports and moves them from the blocking state (if they were in the blocking state), through the listening and learning states, and into the forwarding state.

Figure 22-5 shows an example topology with no link failures. Switch A, the root switch, connects directly to Switch B over link L1 and to Switch C over link L2. The Layer 2 interface on Switch C that connects directly to Switch B is in the blocking state.

Figure 22-5 BackboneFast Example Before Indirect Link Failure

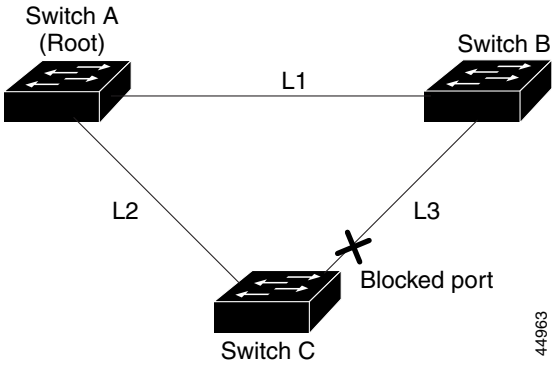

If link L1 fails as shown in Figure 22-6, Switch C cannot detect this failure because it is not connected directly to link L1. However, because Switch B is directly connected to the root switch over L1, it detects the failure, elects itself the root, and begins sending BPDUs to Switch C, identifying itself as the root. When Switch C receives the inferior BPDUs from Switch B, Switch C assumes that an indirect failure has occurred. At that point, BackboneFast allows the blocked interface on Switch C to move immediately to the listening state without waiting for the maximum aging time for the interface to expire. BackboneFast then transitions the Layer 2 interface on Switch C to the forwarding state, providing a path from Switch B to Switch A. The root-switch election takes approximately 30 seconds, twice the Forward Delay time if the default Forward Delay time of 15 seconds is set. Figure 22-6 shows how BackboneFast reconfigures the topology to account for the failure of link L1.

Figure 22-6 BackboneFast Example After Indirect Link Failure

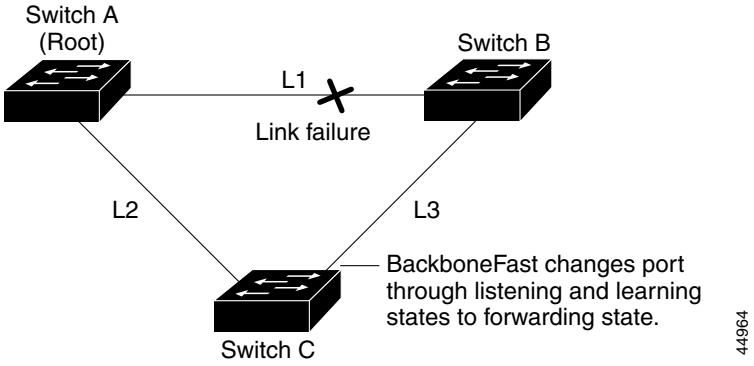

If a new switch is introduced into a shared-medium topology as shown in Figure 22-7, BackboneFast is not activated because the inferior BPDUs did not come from the recognized designated switch (Switch B). The new switch begins sending inferior BPDUs that indicate it is the root switch. However, the other switches ignore these inferior BPDUs, and the new switch learns that Switch B is the designated switch to Switch A, the root switch.

Figure 22-7 Adding a Switch in a Shared-Medium Topology

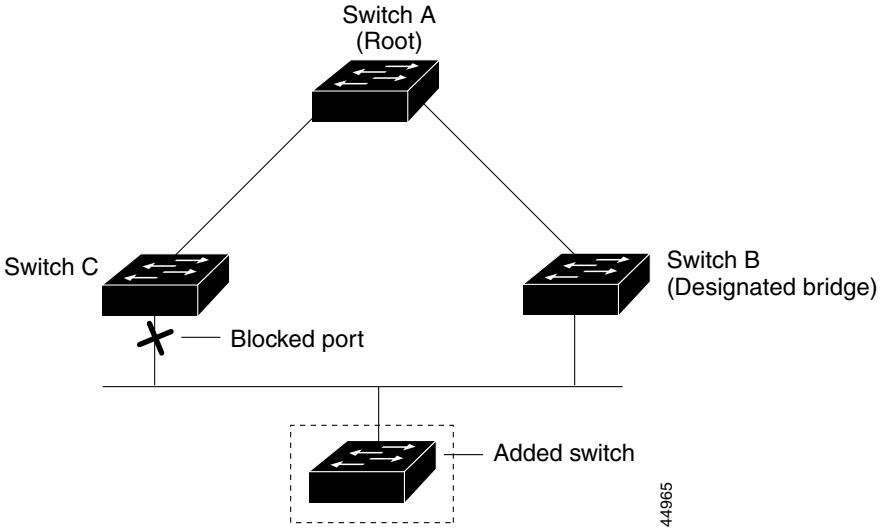

EtherChannel Guard

You can use EtherChannel guard to detect an EtherChannel misconfiguration between the switch and a connected device. A misconfiguration can occur if the switch interfaces are configured in an EtherChannel, but the interfaces on the other device are not. A misconfiguration can also occur if the channel parameters are not the same at both ends of the EtherChannel. For EtherChannel configuration guidelines, see the "EtherChannel Configuration Guidelines" section on page 40-10.

If the switch detects a misconfiguration on the other device, EtherChannel guard places the switch interfaces in the error-disabled state, and displays an error message.

You can enable this feature by using the **spanning-tree etherchannel guard misconfig** global configuration command.

Root Guard

The Layer 2 network of a service provider (SP) can include many connections to switches that are not owned by the SP. In such a topology, the spanning tree can reconfigure itself and select a *customer switch* as the root switch, as shown in Figure 22-8. You can avoid this situation by enabling root guard on SP switch interfaces that connect to switches in your customer's network. If spanning-tree calculations cause an interface in the customer network to be selected as the root port, root guard then places the interface in the root-inconsistent (blocked) state to prevent the customer's switch from becoming the root switch or being in the path to the root.

If a switch outside the SP network becomes the root switch, the interface is blocked (root-inconsistent state), and spanning tree selects a new root switch. The customer's switch does not become the root switch and is not in the path to the root.

If the switch is operating in multiple spanning-tree (MST) mode, root guard forces the interface to be a designated port. If a boundary port is blocked in an internal spanning-tree (IST) instance because of root guard, the interface also is blocked in all MST instances. A boundary port is an interface that connects to a LAN, the designated switch of which is either an IEEE 802.1D switch or a switch with a different MST region configuration.

Root guard enabled on an interface applies to all the VLANs to which the interface belongs. VLANs can be grouped and mapped to an MST instance.

You can enable this feature by using the **spanning-tree guard root** interface configuration command.

Caution Misuse of the root guard feature can cause a loss of connectivity.

Figure 22-8 Root Guard in a Service-Provider Network

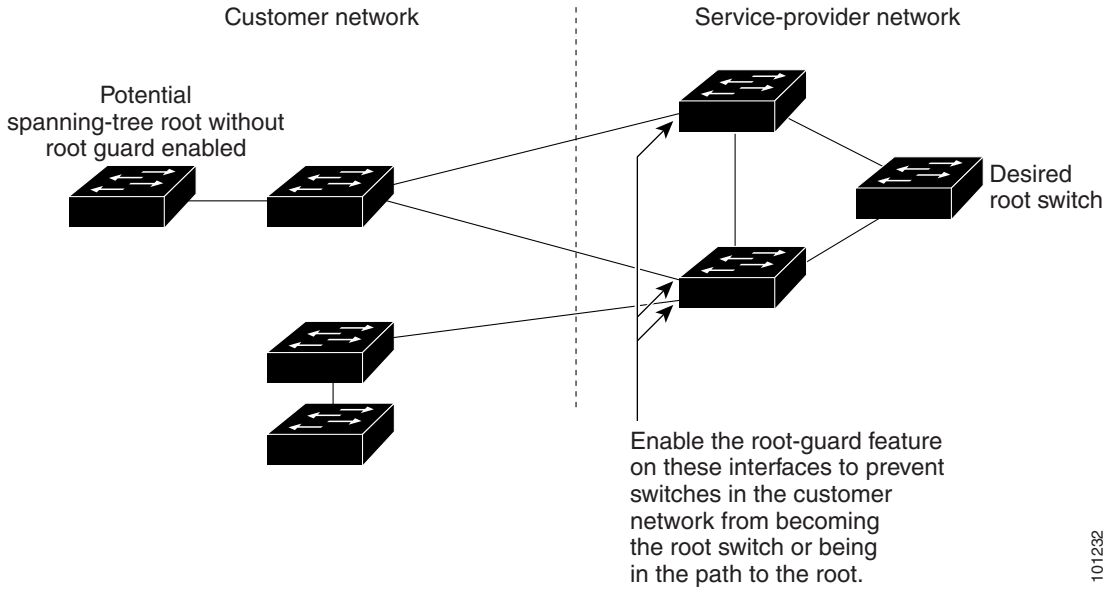

Loop Guard

You can use loop guard to prevent alternate or root ports from becoming designated ports because of a failure that leads to a unidirectional link. This feature is most effective when it is enabled on the entire switched network. Loop guard prevents alternate and root ports from becoming designated ports, and spanning tree does not send BPDUs on root or alternate ports.

You can enable this feature by using the **spanning-tree loopguard default** global configuration command.

When the switch is operating in PVST+ or rapid-PVST+ mode, loop guard prevents alternate and root ports from becoming designated ports, and spanning tree does not send BPDUs on root or alternate ports.

When the switch is operating in MST mode, BPDUs are not sent on nonboundary ports only if the interface is blocked by loop guard in all MST instances. On a boundary port, loop guard blocks the interface in all MST instances.

Default Optional Spanning-Tree Settings

Table 22-1 Default Optional Spanning-Tree Settings

Feature	Default Setting
PortFast, BPDU filtering, BPDU guard	Globally disabled (unless they are individually configured per interface).
UplinkFast	Globally disabled.
BackboneFast	Globally disabled.
EtherChannel guard	Globally enabled.
Root guard	Disabled on all interfaces.
Loop guard	Disabled on all interfaces.

How to Configure the Optional Spanning-Tree Features

Enabling Optional SPT Features

Before You Begin

- Make sure that there are no loops in the network between the trunk port and the workstation or server before you enable PortFast on a trunk port.
- Use PortFast *only* when connecting a single end station to an access or trunk port. Enabling this feature on an interface connected to a switch or hub could prevent spanning tree from detecting and disabling loops in your network, which could cause broadcast storms and address-learning problems.
- An interface with the PortFast feature enabled is moved directly to the spanning-tree forwarding state without waiting for the standard forward-time delay.
- You cannot enable both loop guard and root guard at the same time.
- When you enable UplinkFast, it affects all VLANs on the switch. You cannot configure UplinkFast on an individual VLAN.
- If you enable the voice VLAN feature, the PortFast feature is automatically enabled. When you disable voice VLAN, the PortFast feature is not automatically disabled.

	Command	Purpose
Step 1	**show spanning-tree active** or **show spanning-tree mst**	Verifies which interfaces are alternate or root ports.
Step 2	**configure terminal**	Enters global configuration mode.
Step 3	**spanning-tree loopguard default**	Enables loop guard. By default, loop guard is disabled.

	Command	Purpose
Step 4	**spanning-tree portfast bpduguard default**	Enables BPDU guard. By default, BPDU guard is disabled.
Step 5	**spanning-tree portfast bpdufilter default**	Enables BPDU filtering. By default, BPDU filtering is disabled.
Step 6	**spanning-tree uplinkfast** [**max-update-rate** *pkts-per-second*]	Enables UplinkFast. (Optional) *pkts-per-second*—The range is 0 to 32000 packets per second; the default is 150. If you set the rate to 0, station-learning frames are not generated, and the spanning-tree topology converges more slowly after a loss of connectivity.
Step 7	**spanning-tree backbonefast**	Enables BackboneFast.
Step 8	**spanning-tree etherchannel guard misconfig**	Enables EtherChannel guard.
Step 9	**interface** *interface-id*	Specifies an interface to configure, and enters interface configuration mode.
Step 10	**spanning-tree portfast** [**trunk**]	Enables PortFast on an access port connected to a single workstation or server. By specifying the **trunk** keyword, you can enable PortFast on a trunk port. **Note** To enable PortFast on trunk ports, you must use the **spanning-tree portfast trunk** interface configuration command. The **spanning-tree portfast** command will not work on trunk ports. By default, PortFast is disabled on all interfaces.
Step 11	**spanning-tree guard root**	Enables root guard on the interface. By default, root guard is disabled on all interfaces.
Step 12	**end**	Returns to privileged EXEC mode.

Maintaining and Monitoring Optional Spanning-Tree Features

Command	Purpose
show spanning-tree active	Displays spanning-tree information on active interfaces only.
show spanning-tree detail	Displays a detailed summary of interface information.
show spanning-tree interface *interface-id*	Displays spanning-tree information for the specified interface.
show spanning-tree mst interface *interface-id*	Displays MST information for the specified interface.
show spanning-tree summary [**totals**]	Displays a summary of interface states or displays the total lines of the spanning-tree state section.
show interfaces status err-disabled	Displays which switch ports are disabled because of an EtherChannel misconfiguration.

Command	Purpose
show etherchannel summary	Displays the EtherChannel configuration. Useful to use on the remote device after switch ports are disabled.
[**no**] **shutdown**	Disables the interface. The **no** option reenables the interface.

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
VLAN configuration	*Chapter 17, "Configuring VLANs"*
Voice VLAN configuration	*Chapter 19, "Configuring Voice VLAN"*
PVST+ and rapid PVST+ configuratio	*Chapter 20, "Configuring STP"*
Multiple Spanning Tree Protocol configuration	*Chapter 21, "Configuring MSTP"*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

CHAPTER 23

Configuring Resilient Ethernet Protocol

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for REP

You can configure all of these features when your switch is running the Per VLAN Spanning-Tree Plus (PVST+). You can configure only the noted features when your switch is running the Multiple Spanning Tree Protocol (MSTP) or the Rapid PVST+(RPVST+) protocol.

Restrictions for REP

You can configure the UplinkFast or the BackboneFast feature for Rapid PVST+ or for the MSTP, but the feature remains disabled (inactive) until you change the spanning-tree mode to PVST+.

Information About Configuring REP

REP

Resilient Ethernet Protocol (REP) is a Cisco proprietary protocol that provides an alternative to Spanning Tree Protocol (STP) to control network loops, handle link failures, and improve convergence time. REP controls a group of ports connected in a segment, ensures that the segment does not create any bridging loops, and responds to link failures within the segment. REP provides a basis for constructing more complex networks and supports VLAN load balancing.

One REP segment is a chain of ports connected to each other and configured with a segment ID. Each segment consists of standard (non-edge) segment ports and two user-configured edge ports. A switch can have no more than two ports that belong to the same segment, and each segment port can have only one external neighbor. A segment can go through a shared medium, but on any link only two ports can belong to the same segment. REP is supported only on Layer 2 trunk interfaces.

Figure 23-1 shows an example of a segment consisting of six ports spread across four switches. Ports E1 and E2 are configured as edge ports. When all ports are operational (as in the segment on the left), a single port is blocked, shown by the diagonal line. When there is a failure in the network, as shown in the diagram on the right, the blocked port returns to the forwarding state to minimize network disruption.

Figure 23-1 REP Open Segments

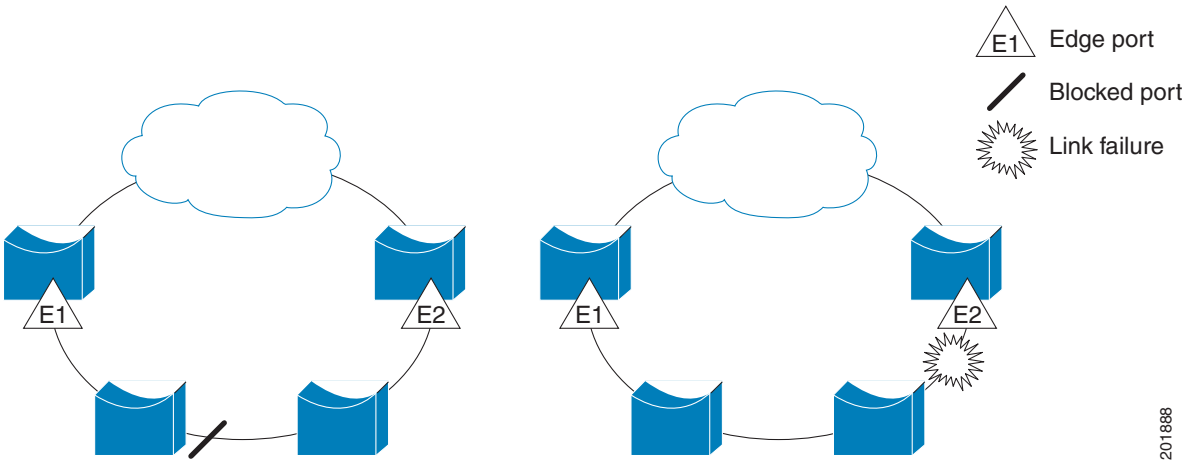

The segment shown in Figure 23-1 is an open segment; there is no connectivity between the two edge ports. The REP segment cannot cause a bridging loop and it is safe to connect the segment edges to any network. All hosts connected to switches inside the segment have two possible connections to the rest of the network through the edge ports, but only one connection is accessible at any time. If a failure causes a host to be unable to access its usual gateway, REP unblocks all ports to ensure that connectivity is available through the other gateway.

The segment shown in Figure 23-2, with both edge ports located on the same switch, is a ring segment. In this configuration, there is connectivity between the edge ports through the segment. With this configuration, you can create a redundant connection between any two switches in the segment.

Figure 23-2 REP Ring Segment

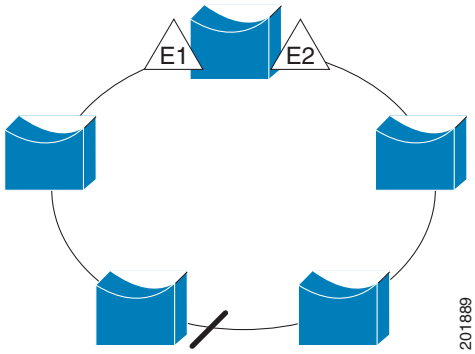

REP segments have these characteristics:

- If all ports in the segment are operational, one port (referred to as the *alternate* port) is in the blocked state for each VLAN. If VLAN load balancing is configured, two ports in the segment control the blocked state of VLANs.
- If one or more ports in a segment is not operational, causing a link failure, all ports forward traffic on all VLANs to ensure connectivity.
- In case of a link failure, the alternate ports are unblocked as quickly as possible. When the failed link comes back up, a logically blocked port per VLAN is selected with minimal disruption to the network.

You can construct almost any type of network based on REP segments. REP also supports VLAN load-balancing, controlled by the primary edge port but occurring at any port in the segment.

In access ring topologies, the neighboring switch might not support REP, as shown in Figure 23-3. In this case, you can configure the non-REP facing ports (E1 and E2) as edge no-neighbor ports. These ports inherit all properties of edge ports, and you can configure them the same as any edge port, including configuring them to send STP or REP topology change notices to the aggregation switch. In this case the STP topology change notice (TCN) that is sent is a multiple spanning-tree (MST) STP message.

Figure 23-3 Edge No-Neighbor Ports

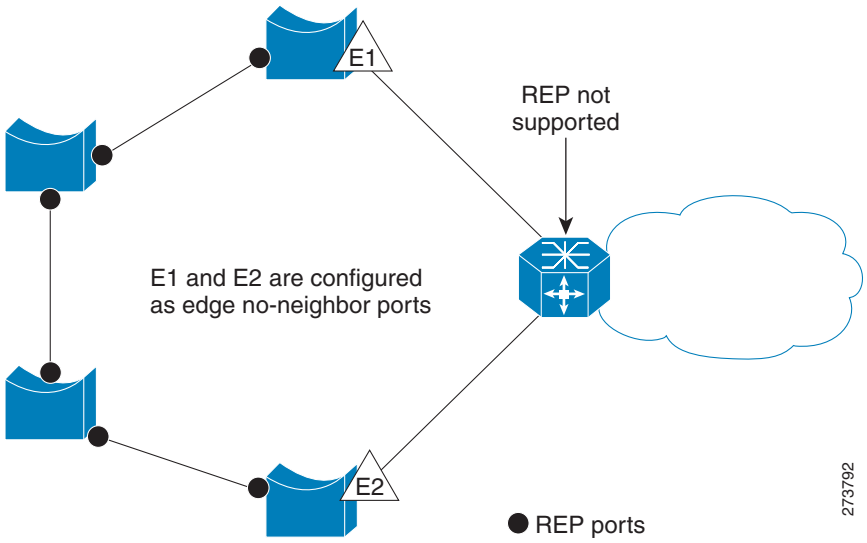

REP has these limitations:

- You must configure each segment port; an incorrect configuration can cause forwarding loops in the networks.
- REP can manage only a single failed port within the segment; multiple port failures within the REP segment cause loss of network connectivity.
- You should configure REP only in networks with redundancy. Configuring REP in a network without redundancy causes loss of connectivity.

Link Integrity

REP does not use an end-to-end polling mechanism between edge ports to verify link integrity. It implements local link failure detection. The REP Link Status Layer (LSL) detects its REP-aware neighbor and establishes connectivity within the segment. All VLANs are blocked on an interface until it detects the neighbor. After the neighbor is identified, REP determines which neighbor port should become the alternate port and which ports should forward traffic.

Each port in a segment has a unique port ID. The port ID format is similar to that used by the spanning tree algorithm: a port number (unique on the bridge), associated to a MAC address (unique in the network). When a segment port is coming up, its LSL starts sending packets that include the segment ID and the port ID. The port is declared as operational after it performs a three-way handshake with a neighbor in the same segment.

A segment port does not become operational if:

- No neighbor has the same segment ID.
- More than one neighbor has the same segment ID.
- The neighbor does not acknowledge the local port as a peer.

Each port creates an adjacency with its immediate neighbor. Once the neighbor adjacencies are created, the ports negotiate to determine one blocked port for the segment, the alternate port. All other ports become unblocked. By default, REP packets are sent to a BPDU class MAC address. The packets can also be sent to the Cisco multicast address, which is used only to send blocked port advertisement (BPA) messages when there is a failure in the segment. The packets are dropped by devices not running REP.

Fast Convergence

Because REP runs on a physical link basis and not a per-VLAN basis, only one hello message is required for all VLANs, reducing the load on the protocol. We recommend that you create VLANs consistently on all switches in a given segment and configure the same allowed VLANs on the REP trunk ports. To avoid the delay introduced by relaying messages in software, REP also allows some packets to be flooded to a regular multicast address. These messages operate at the hardware flood layer (HFL) and are flooded to the whole network, not just the REP segment. Switches that do not belong to the segment treat them as data traffic. You can control flooding of these messages by configuring a dedicated administrative VLAN for the whole domain.

The estimated convergence recovery time on fiber interfaces is less than 200 ms for the local segment with 200 VLANs configured. Convergence for VLAN load balancing is 300 ms or less.

VLAN Load Balancing

One edge port in the REP segment acts as the primary edge port; the other as the secondary edge port. It is the primary edge port that always participates in VLAN load balancing in the segment. REP VLAN balancing is achieved by blocking some VLANs at a configured alternate port and all other VLANs at the primary edge port. When you configure VLAN load balancing, you can specify the alternate port in one of three ways:

- By entering the port ID of the interface. To identify the port ID of a port in the segment, enter the **show interface rep detail** interface configuration command for the port.

- By entering the neighbor offset number of a port in the segment, which identifies the downstream neighbor port of an edge port. The neighbor offset number range is –256 to +256; a value of 0 is invalid. The primary edge port has an offset number of 1; positive numbers above 1 identify downstream neighbors of the primary edge port. Negative numbers indicate the secondary edge port (offset number -1) and its downstream neighbors.

Note You configure offset numbers on the primary edge port by identifying a port's downstream position from the primary (or secondary) edge port. You would never enter an offset value of 1 because that is the offset number of the primary edge port itself.

Figure 23-4 shows neighbor offset numbers for a segment where E1 is the primary edge port and E2 is the secondary edge port. The red numbers inside the ring are numbers offset from the primary edge port; the black numbers outside of the ring show the offset numbers from the secondary edge port. Note that you can identify all ports (except the primary edge port) by either a positive offset number (downstream position from the primary edge port) or a negative offset number (downstream position from the secondary edge port). If E2 became the primary edge port, its offset number would then be 1 and E1 would be -1.

- By entering the **preferred** keyword to select the port that you previously configured as the preferred alternate port with the **rep segment** *segment-id* **preferred** interface configuration command.

Figure 23-4 Neighbor Offset Numbers in a Segment

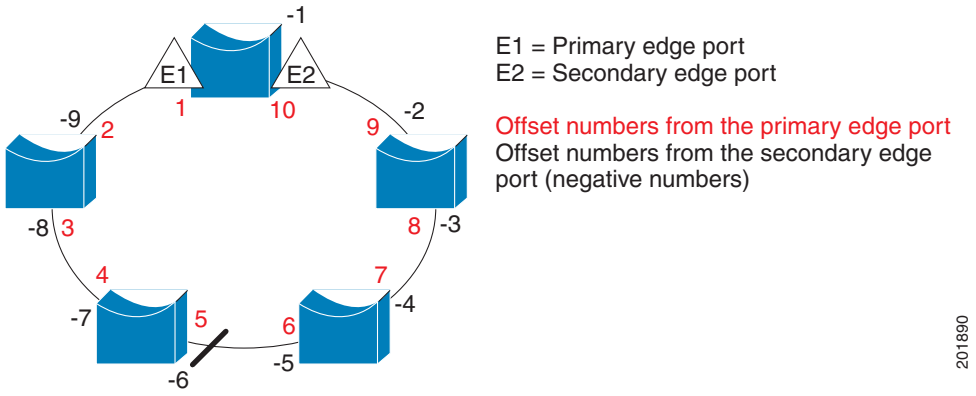

When the REP segment is complete, all VLANs are blocked. When you configure VLAN load balancing, you must also configure triggers in one of two ways:

- Manually trigger VLAN load balancing at any time by entering the **rep preempt segment** *segment-id* privileged EXEC command on the switch that has the primary edge port.
- Configure a preempt delay time by entering the **rep preempt delay** *seconds* interface configuration command. After a link failure and recovery, VLAN load balancing begins after the configured preemption time period elapses. Note that the delay timer restarts if another port fails before the time has elapsed.

Note When VLAN load balancing is configured, it does not start working until triggered by either manual intervention or a link failure and recovery.

When VLAN load balancing is triggered, the primary edge port sends out a message to alert all interfaces in the segment about the preemption. When the secondary port receives the message, it is reflected into the network to notify the alternate port to block the set of VLANs specified in the message and to notify the primary edge port to block the remaining VLANs.

You can also configure a particular port in the segment to block all VLANs. Only the primary edge port initiates VLAN load balancing, which is not possible if the segment is not terminated by an edge port on each end. The primary edge port determines the local VLAN load balancing configuration.

Reconfigure the primary edge port to reconfigure load balancing. When you change the load balancing configuration, the primary edge port again waits for the **rep preempt segment** command or for the configured preempt delay period after a port failure and recovery before executing the new configuration. If you change an edge port to a regular segment port, the existing VLAN load balancing status does not change. Configuring a new edge port might cause a new topology configuration.

Spanning Tree Interaction

REP does not interact with STP or with the FlexLink feature, but can coexist with both. A port that belongs to a segment is removed from spanning tree control and STP BPDUs are not accepted or sent from segment ports. Therefore, STP cannot run on a segment.

To migrate from an STP ring configuration to REP segment configuration, begin by configuring a single port in the ring as part of the segment and continue by configuring contiguous ports to minimize the number of segments. Each segment always contains a blocked port, so multiple segments means multiple blocked ports and a potential loss of connectivity. When the segment has been configured in both directions up to the location of the edge ports, you then configure the edge ports.

REP Ports

Ports in REP segments are Failed, Open, or Alternate.

- A port configured as a regular segment port starts as a failed port.
- After the neighbor adjacencies are determined, the port transitions to alternate port state, blocking all VLANs on the interface. Blocked port negotiations occur and when the segment settles, one blocked port remains in the alternate role and all other ports become open ports.
- When a failure occurs in a link, all ports move to the failed state. When the alternate port receives the failure notification, it changes to the open state, forwarding all VLANs.

A regular segment port converted to an edge port, or an edge port converted to a regular segment port, does not always result in a topology change. If you convert an edge port into a regular segment port, VLAN load balancing is not implemented unless it has been configured. For VLAN load balancing, you must configure two edge ports in the segment.

A segment port that is reconfigured as a spanning tree port restarts according the spanning tree configuration. By default, this is a designated blocking port. If PortFast is configured or if STP is disabled, the port goes into the forwarding state.

REP Segments

A segment is a collection of ports connected one to the other in a chain and configured with a segment ID. To configure REP segments, you configure the REP administrative VLAN (or use the default VLAN 1) and then add the ports to the segment using interface configuration mode. You should configure two edge ports in the segment, with one of them the primary edge port and the other by default the secondary edge port. A segment has only one primary edge port. If you configure two ports in a segment as the primary edge port, for example, ports on different switches, the REP selects one of them to serve as the segment primary edge port. You can also optionally configure where to send segment topology change notices (STCNs) and VLAN load balancing.

Default REP Configuration

REP is disabled on all interfaces. When enabled, the interface is a regular segment port unless it is configured as an edge port.

When REP is enabled, the sending of segment topology change notices (STCNs) is disabled, all VLANs are blocked, and the administrative VLAN is VLAN 1.

When VLAN load balancing is enabled, the default is manual preemption with the delay timer disabled. If VLAN load balancing is not configured, the default after manual preemption is to block all VLANs at the primary edge port.

REP Configuration Guidelines

Follow these guidelines when configuring REP:

- We recommend that you begin by configuring one port and then configure the contiguous ports to minimize the number of segments and the number of blocked ports.
- If more than two ports in a segment fail when no external neighbors are configured, one port changes to a forwarding state for the data path to help maintain connectivity during configuration. In the show rep interface privileged EXEC command output, the Port Role for this port shows as *Fail Logical Open*; the Port Role for the other failed port shows as *Fail No Ext Neighbor.* When the external neighbors for the failed ports are configured, the ports go through the alternate port state transitions and eventually go to an open state or remain as the alternate port, based on the alternate port election mechanism.
- REP ports must be Layer 2 trunk ports.
- Be careful when configuring REP through a Telnet connection. Because REP blocks all VLANs until another REP interface sends a message to unblock it, you might lose connectivity to the switch if you enable REP in a Telnet session that accesses the switch through the same interface.
- You cannot run REP and STP or REP and Flex Links on the same segment or interface.
- If you connect an STP network to the REP segment, be sure that the connection is at the segment edge. An STP connection that is not at the edge could cause a bridging loop because STP does not run on REP segments. All STP BPDUs are dropped at REP interfaces.
- You must configure all trunk ports in the segment with the same set of allowed VLANs, or a misconfiguration occurs.

- REP ports follow these rules:
  - There is no limit to the number of REP ports on a switch; however, only two ports on a switch can belong to the same REP segment.
  - If only one port on a switch is configured in a segment, the port should be an edge port.
  - If two ports on a switch belong to the same segment, they must be both edge ports, both regular segment ports, or one regular port and one edge no-neighbor port. An edge port and regular segment port on a switch cannot belong to the same segment.
  - If two ports on a switch belong to the same segment and one is configured as an edge port and one as a regular segment port (a misconfiguration), the edge port is treated as a regular segment port.
- REP interfaces come up in a blocked state and remains in a blocked state until notified that it is safe to unblock. You need to be aware of this to avoid sudden connection losses.
- REP sends all LSL PDUs in untagged frames on the native VLAN. The BPA message sent to the Cisco multicast address is sent on the administration VLAN, which is VLAN 1 by default.
- You can configure how long a REP interface remains up without receiving a hello from a neighbor. You can use the **rep lsl-age-timer** *value* interface configuration command to set the time from 120 ms to 10000 ms. The LSL hello timer is then set to the age-timer value divided by 3. In normal operation, three LSL hellos are sent before the age timer on the peer switch expires and checks for hello messages.
  - In Cisco IOS Release 12.2(52)SE, the LSL age-timer range changed from 3000 to 10000 ms in 500-ms increments to 120 to 10000 ms in 40-ms increments. If the REP neighbor device is not running Cisco IOS release 12.2(52)SE or later, do not configure a timer value less than 3000 ms. Configuring a value less than 3000 ms causes the port to shut down because the neighbor switch does not respond within the requested time period.
  - EtherChannel port channel interfaces do not support LSL age-timer values less than 1000 ms. If you try to configure a value less than 1000 ms on a port channel, you receive an error message and the command is rejected.
- When configuring the REP LSL age timer, make sure that both ends of the link have the same time value configured. Configuring different values on ports at each end of the link results in a REP link flap.
- REP ports cannot be configured as one of these port types:
  - SPAN destination port
  - Tunnel port
  - Access port
- REP is supported on EtherChannels, but not on an individual port that belongs to an EtherChannel.
- There is a maximum of 64 REP segments per switch.

REP Administrative VLAN

To avoid the delay introduced by relaying messages in software for link-failure or VLAN-blocking notification during load balancing, REP floods packets at the hardware flood layer (HFL) to a regular multicast address. These messages are flooded to the whole network, not just the REP segment. You can control flooding of these messages by configuring an administrative VLAN for the whole domain.

Follow these guidelines when configuring the REP administrative VLAN:

- If you do not configure an administrative VLAN, the default is VLAN 1.
- There can be only one administrative VLAN on a switch and on a segment. However, this is not enforced by software.
- The administrative VLAN cannot be the RSPAN VLAN.

How to Configure REP

Configuring the REP Administrative VLAN

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**rep admin vlan** *vlan-id*	Specifies the administrative VLAN. The range is 2 to 4096. The default is VLAN 1. To set the admin VLAN to 1, enter the **no rep admin vlan** global configuration command.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring REP Interfaces

Before You Begin

For REP operation, you need to enable it on each segment interface and identify the segment ID. This step is required and must be done before other REP configuration. You must also configure a primary and secondary edge port on each segment. All other steps are optional.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface, and enters interface configuration mode. The interface can be a physical Layer 2 interface or a port channel (logical interface). The port-channel range is 1 to 48.
Step 3	**switchport mode trunk**	Configures the interface as a Layer 2 trunk port.

	Command	Purpose
Step 4	**rep segment** *segment-id* [**edge** [**no-neighbor**] [**primary**]] [**preferred**]	Enables REP on the interface, and identifies a segment number. The segment ID range is from 1 to 1024. These optional keywords are available: **Note** You must configure two edge ports, including one primary edge port for each segment. • **edge**—Configures the port as an edge port. Entering **edge** without the **primary** keyword configures the port as the secondary edge port. Each segment has only two edge ports. • (Optional) **primary**— Configures the port as the primary edge port, the port on which you can configure VLAN load balancing. • (Optional) **no-neighbor—**Configures a port with no external REP neighbors as an edge port. The port inherits all properties of edge ports, and you can configure them the same as any edge port. **Note** Although each segment can have only one primary edge port, if you configure edge ports on two different switches and enter the **primary** keyword on both switches, the configuration is allowed. However, REP selects only one of these ports as the segment primary edge port. You can identify the primary edge port for a segment by entering the **show rep topology** privileged EXEC command. • (Optional) **preferred**—Indicates that the port is the preferred alternate port or the preferred port for VLAN load balancing. **Note** Configuring a port as preferred does not guarantee that it becomes the alternate port; it merely gives it a slight edge among equal contenders. The alternate port is usually a previously failed port.
Step 5	**rep stcn** {**interface** *interface-id* | **segment** *id-list* | **stp**}	(Optional) Configures the edge port to send segment topology change notices (STCNs). • **interface** *interface-id*—Designates a physical interface or port channel to receive STCNs. • **segment** *id-list*—Identifies one or more segments to receive STCNs. The range is 1 to 1024. • **stp—**Sends STCNs to STP networks.

	Command	Purpose
Step 6	**rep block port** {**id** *port-id* | *neighbor_offset* | **preferred**} **vlan** {*vlan-list* | **all**}	(Optional) Configures VLAN load balancing on the primary edge port, identify the REP alternate port in one of three ways, and configure the VLANs to be blocked on the alternate port. • **id** *port-id*—Identifies the alternate port by port ID. The port ID is automatically generated for each port in the segment. You can view interface port IDs by entering the **show interface** *interface-id* **rep** [**detail**] privileged EXEC command. • *neighbor_offset* number—Identifies the alternate port as a downstream neighbor from an edge port. The range is from –256 to 256, with negative numbers indicating the downstream neighbor from the secondary edge port. A value of **0** is invalid. Enters **-1** to identify the secondary edge port as the alternate port. See Figure 23-4 on page 23-5 for an example of neighbor offset numbering. **Note** Because you enter this command at the primary edge port (offset number 1), you would never enter an offset value of 1 to identify an alternate port. • **preferred**—Selects the regular segment port previously identified as the preferred alternate port for VLAN load balancing. • **vlan** *vlan-list*—Blocks one VLAN or a range of VLANs. • **vlan all**—Blocks all VLANs. **Note** Enter this command only on the REP primary edge port.
Step 7	**rep preempt delay** *seconds*	(Optional) You must enter this command and configure a preempt time delay if you want VLAN load balancing to automatically trigger after a link failure and recovery. The time delay range is 15 to 300 seconds. The default is manual preemption with no time delay. **Note** Enter this command only on the REP primary edge port.
Step 8	**rep lsl-age-timer** *value*	(Optional) Configures a time (in milliseconds) for which the REP interface remains up without receiving a hello from a neighbor. The range is from 120 to 10000 ms in 40-ms increments. The default is 5000 ms (5 seconds). **Note** If the neighbor device is not running Cisco IOS Release 12.2(52)SE or later, it only accepts values from 3000 to 10000 ms in 500-ms intervals. EtherChannel port channel interfaces do not support LSL age-timer values less than 1000 ms.
Step 9	**end**	Returns to privileged EXEC mode.

Setting Manual Preemption for VLAN Load Balancing

Before You Begin

If you do not enter the **rep preempt delay** *seconds* interface configuration command on the primary edge port to configure a preemption time delay, the default is to manually trigger VLAN load balancing on the segment. Be sure that all other segment configuration has been completed before manually preempting VLAN load balancing. When you enter the **rep preempt segment** *segment-id* command, a confirmation message appears before the command is executed because preemption can cause network disruption.

	Command	Purpose
Step 1	**rep preempt segment** *segment-id*	Manually triggers VLAN load balancing on the segment. You will need to confirm the command before it is executed.
Step 2	**show rep topology**	Displays REP topology information.

Configuring SNMP Traps for REP

You can configure the switch to send REP-specific traps to notify the SNMP server of link operational status changes and port role changes.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**snmp mib rep trap-rate** *value*	Enables the switch to send REP traps, and sets the number of traps sent per second. The range is from 0 to 1000. The default is 0 (no limit imposed; a trap is sent at every occurrence).
Step 3	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining REP

Command	Purpose
show interface [*interface-id*] **rep** [**detail**]	Displays REP configuration and status for an interface or for all interfaces.
show rep topology [**segment** *segment_id*] [**archive**] [**detail**]	Displays REP topology information for a segment or for all segments, including the primary and secondary edge ports in the segment.
copy running-config startup config	Saves your entries in the switch startup configuration file.

Configuration Examples for Configuring REP

Configuring the Administrative VLAN: Example

This example shows how to configure the administrative VLAN as VLAN 100 and verify the configuration by entering the **show interface rep detail** command on one of the REP interfaces:

```
Switch# configure terminal
Switch (conf)# rep admin vlan 100
Switch (conf-if)# end
Switch# show interface gigabitethernet1/1 rep detail
GigabitEthernet1/1 REP enabled
Segment-id: 2 (Edge)
PortID: 00010019E7144680
Preferred flag: No
Operational Link Status: TWO_WAY
Current Key: 0002001121A2D5800E4D
Port Role: Open
Blocked Vlan: <empty>
Admin-vlan: 100
Preempt Delay Timer: disabled
LSL Ageout Timer: 5000 ms
Configured Load-balancing Block Port: none
Configured Load-balancing Block VLAN: none
STCN Propagate to: none
LSL PDU rx: 3322, tx: 1722
HFL PDU rx: 32, tx: 5
BPA TLV rx: 16849, tx: 508
BPA (STCN, LSL) TLV rx: 0, tx: 0
BPA (STCN, HFL) TLV rx: 0, tx: 0
EPA-ELECTION TLV rx: 118, tx: 118
EPA-COMMAND TLV rx: 0, tx: 0
EPA-INFO TLV rx: 4214, tx: 4190
```

Configuring a Primary Edge Port: Examples

This example shows how to configure an interface as the primary edge port for segment 1, to send STCNs to segments 2 through 5, and to configure the alternate port as the port with port ID 0009001818D68700 to block all VLANs after a preemption delay of 60 seconds after a segment port failure and recovery. The interface is configured to remain up for 6000 milliseconds without receiving a hello from a neighbor.

```
Switch# configure terminal
Switch (conf)# interface gigabitethernet1/1
Switch (conf-if)# rep segment 1 edge primary
Switch (conf-if)# rep stcn segment 2-5
Switch (conf-if)# rep block port 0009001818D68700 vlan all
Switch (conf-if)# rep preempt delay 60
Switch (conf-if)# rep lsl-age-timer 6000
Switch (conf-if)# end
```

This example shows how to configure an interface as the primary edge port when the interface has no external REP neighbor:

```
Switch# configure terminal
Switch (conf)# interface gigabitethernet1/1
Switch (conf-if)# rep segment 1 edge no-neighbor primary
Switch (conf-if)# rep stcn segment 2-5
Switch (conf-if)# rep block port 0009001818D68700 vlan all
```

```
Switch (conf-if)# rep preempt delay 60
Switch (conf-if)# rep lsl-age-timer 6000
```

Configuring VLAN Blocking: Example

This example shows how to configure the VLAN blocking configuration shown in Figure 23-5. The alternate port is the neighbor with neighbor offset number 4. After manual preemption, VLANs 100 to 200 are blocked at this port, and all other VLANs are blocked at the primary edge port E1 (Gigabit Ethernet port 1/0/1).

```
Switch# configure terminal
Switch (conf)# interface gigabitethernet1/1
Switch (conf-if)# rep segment 1 edge primary
Switch (conf-if)# rep block port 4 vlan 100-200
Switch (conf-if)# end
```

Figure 23-5 Example of VLAN Blocking

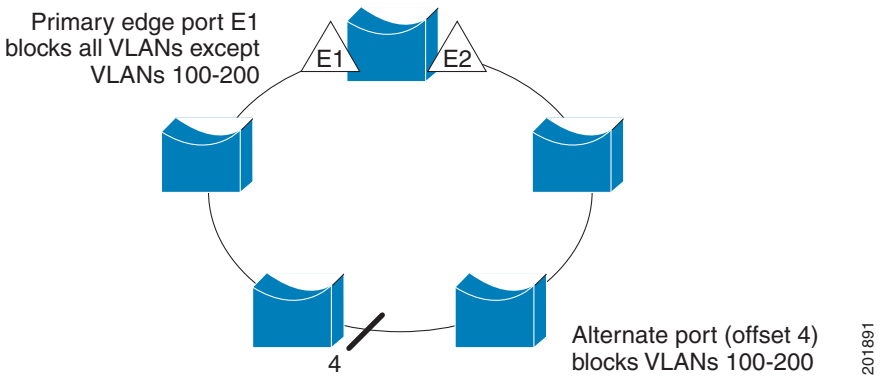

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

CHAPTER 24

Configuring FlexLinks and the MAC Address-Table Move Update

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Restrictions for the FlexLinks and the MAC Address-Table Move Update

- To use this feature, the switch must be running the LAN Base image.

Information About Configuring the FlexLinks and the MAC Address-Table Move Update

FlexLinks

FlexLinks are a pair of a Layer 2 interfaces (switch ports or port channels) where one interface is configured to act as a backup to the other. The feature provides an alternative solution to the Spanning Tree Protocol (STP). Users can disable STP and still retain basic link redundancy. FlexLinks are typically configured in service provider or enterprise networks where customers do not want to run STP on the switch. If the switch is running STP, FlexLinks is not necessary because STP already provides link-level redundancy or backup.

You configure FlexLinks on one Layer 2 interface (the active link) by assigning another Layer 2 interface as the FlexLinks or backup link. When one of the links is up and forwarding traffic, the other link is in standby mode, ready to begin forwarding traffic if the other link shuts down. At any given time, only one

of the interfaces is in the linkup state and forwarding traffic. If the primary link shuts down, the standby link starts forwarding traffic. When the active link comes back up, it goes into standby mode and does not forward traffic. STP is disabled on FlexLinks interfaces.

In Figure 24-1, ports 1 and 2 on switch A are connected to uplink switches B and C. Because they are configured as FlexLinks, only one of the interfaces is forwarding traffic; the other is in standby mode. If port 1 is the active link, it begins forwarding traffic between port 1 and switch B; the link between port 2 (the backup link) and switch C is not forwarding traffic. If port 1 goes down, port 2 comes up and starts forwarding traffic to switch C. When port 1 comes back up, it goes into standby mode and does not forward traffic; port 2 continues forwarding traffic.

You can also choose to configure a preemption mechanism, specifying the preferred port for forwarding traffic. For example, in the example in Figure 24-1, you can configure the FlexLinks pair with preemption mode. In the scenario shown, when port 1 comes back up and has more bandwidth than port 2, port 1 begins forwarding traffic after 60 seconds. Port 2 becomes the standby port. You do this by entering the interface configuration **switchport backup interface preemption mode bandwidth** and **switchport backup interface preemption delay** commands.

Figure 24-1 FlexLinks Configuration Example

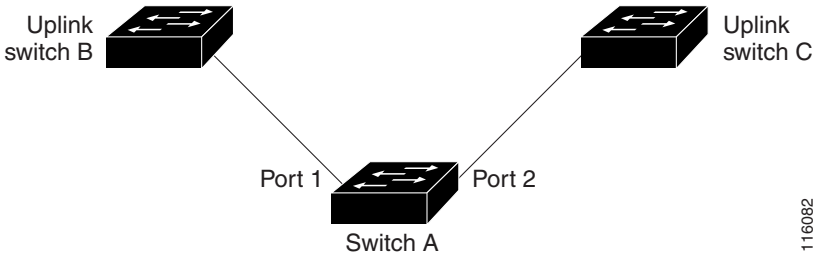

If a primary (forwarding) link goes down, a trap notifies the network management stations. If the standby link goes down, a trap notifies the users.

FlexLinks are supported only on Layer 2 ports and port channels, not on VLANs or on Layer 3 ports.

VLAN FlexLinks Load Balancing and Support

VLAN FlexLinks load-balancing allows you to configure a FlexLinks pair so that both ports simultaneously forward the traffic for some mutually exclusive VLANs. For example, if FlexLinks ports are configured for 1 to100 VLANs, the traffic of the first 50 VLANs can be forwarded on one port and the rest on the other port. If one of the ports fail, the other active port forwards all the traffic. When the failed port comes back up, it resumes forwarding traffic in the preferred VLANs. This way, apart from providing the redundancy, this FlexLinks pair can be used for load balancing. FlexLinks VLAN load balancing does not impose any restrictions on uplink switches.

Figure 24-2 VLAN FlexLinks Load Balancing Configuration Example

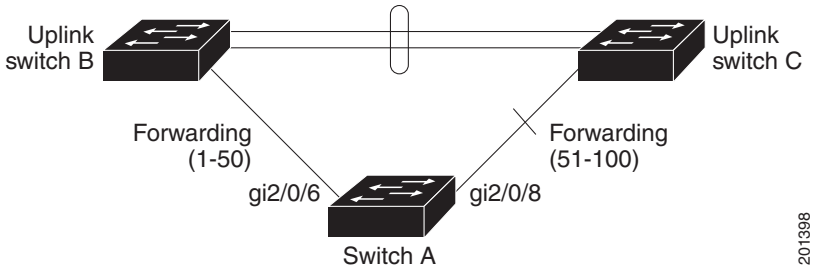

FlexLinks Multicast Fast Convergence

FlexLinks Multicast Fast Convergence reduces the multicast traffic convergence time after a FlexLinks failure.

Learning the Other FlexLinks Port as the mrouter Port

In a typical multicast network, there is a querier for each VLAN. A switch deployed at the edge of a network has one of its FlexLinks ports receiving queries. FlexLinks ports are also always forwarding at any given time.

A port that receives queries is added as an *mrouter* port on the switch. An mrouter port is part of all the multicast groups learned by the switch. After a changeover, queries are received by the other FlexLinks port. The other FlexLinks port is then learned as the mrouter port. After the changeover, multicast traffic flows through the other FlexLinks port. To achieve faster convergence of traffic, both FlexLinks ports are learned as mrouter ports whenever either FlexLinks port is learned as the mrouter port. Both FlexLinks ports are always part of multicast groups.

Though both FlexLinks ports are part of the groups in normal operation mode, all traffic on the backup port is blocked. So the normal multicast data flow is not affected by the addition of the backup port as an mrouter port. When the changeover happens, the backup port is unblocked, allowing the traffic to flow. In this case, the upstream multicast data flows as soon as the backup port is unblocked.

Generating IGMP Reports

When the backup link comes up after the changeover, the upstream new distribution switch does not start forwarding multicast data, because the port on the upstream router, which is connected to the blocked FlexLinks port, is not part of any multicast group. The reports for the multicast groups were not forwarded by the downstream switch because the backup link is blocked. The data does not flow on this port, until it learns the multicast groups, which occurs only after it receives reports.

The reports are sent by hosts when a general query is received, and a general query is sent within 60 seconds in normal scenarios. When the backup link starts forwarding, to achieve faster convergence of multicast data, the downstream switch immediately sends proxy reports for all the learned groups on this port without waiting for a general query.

Leaking IGMP Reports

To achieve multicast traffic convergence with minimal loss, a redundant data path must be set up before the FlexLinks active link goes down. This can be achieved by leaking only IGMP report packets on the FlexLinks backup link. These leaked IGMP report messages are processed by upstream distribution routers, so multicast data traffic gets forwarded to the backup interface. Because all incoming traffic on the backup interface is dropped at the ingress of the access switch, no duplicate multicast traffic is received by the host. When the FlexLinks active link fails, the access switch starts accepting traffic from the backup link immediately. The only disadvantage of this scheme is that it consumes bandwidth on the link between the distribution switches and on the backup link between the distribution and access switches. This feature is disabled by default and can be configured by using the **switchport backup interface** *interface-id* **multicast fast-convergence** command.

When this feature has been enabled at changeover, the switch does not generate the proxy reports on the backup port, which became the forwarding port.

MAC Address-Table Move Update

The MAC address-table move update feature allows the switch to provide rapid bidirectional convergence when a primary (forwarding) link goes down and the standby link begins forwarding traffic.

In Figure 24-3, switch A is an access switch, and ports 1 and 2 on switch A are connected to uplink switches B and D through a FlexLinks pair. Port 1 is forwarding traffic, and port 2 is in the backup state. Traffic from the PC to the server is forwarded from port 1 to port 3. The MAC address of the PC has been learned on port 3 of switch C. Traffic from the server to the PC is forwarded from port 3 to port 1.

If the MAC address-table move update feature is not configured and port 1 goes down, port 2 starts forwarding traffic. However, for a short time, switch C keeps forwarding traffic from the server to the PC through port 3, and the PC does not get the traffic because port 1 is down. If switch C removes the MAC address of the PC on port 3 and relearns it on port 4, traffic can then be forwarded from the server to the PC through port 2.

If the MAC address-table move update feature is configured and enabled on the switches in Figure 24-3 and port 1 goes down, port 2 starts forwarding traffic from the PC to the server. The switch sends a MAC address-table move update packet from port 2. Switch C gets this packet on port 4 and immediately learns the MAC address of the PC on port 4, which reduces the reconvergence time.

You can configure the access switch, switch A, to *send* MAC address-table move update messages. You can also configure the uplink switches B, C, and D to *get* and process the MAC address-table move update messages. When switch C gets a MAC address-table move update message from switch A, switch C learns the MAC address of the PC on port 4. Switch C updates the MAC address table, including the forwarding table entry for the PC.

Switch A does not need to wait for the MAC address-table update. The switch detects a failure on port 1 and immediately starts forwarding server traffic from port 2, the new forwarding port. This change occurs in 100 milliseconds (ms). The PC is directly connected to switch A, and the connection status does not change. Switch A does not need to update the PC entry in the MAC address table.

Figure 24-3 MAC Address-Table Move Update Example

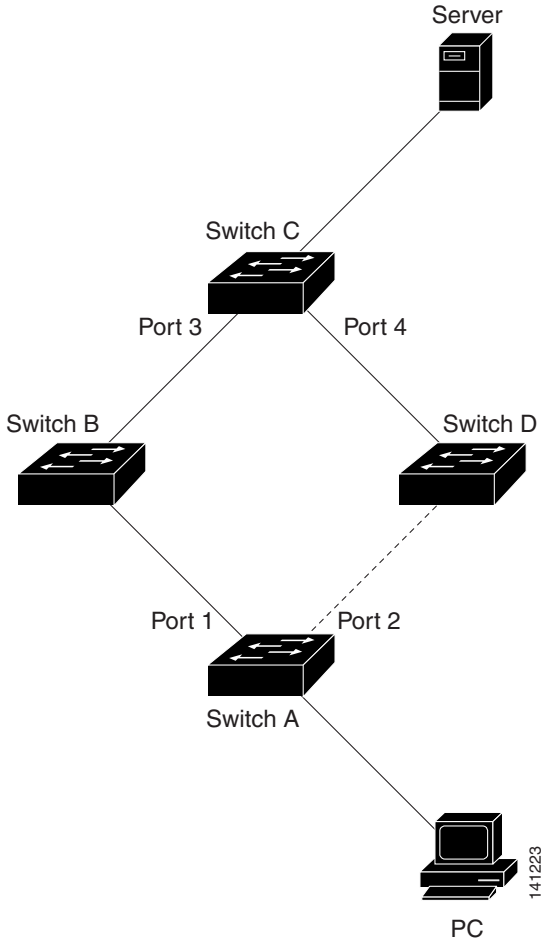

Default Settings for FlexLinks and MAC Address-Table Move Update

Default Settings
FlexLinks is not configured, and there are no backup interfaces defined.
The preemption mode is off.
The preemption delay is 35 seconds.
MAC address-table move update is not configured on the switch.

Configuration Guidelines for FlexLinks and MAC Address-Table Move Update

Follow these guidelines to configure FlexLinks:

- You can configure up to 16 backup links.
- You can configure only one FlexLinks backup link for any active link, and it must be a different interface from the active interface.
- An interface can belong to only one FlexLinks pair. An interface can be a backup link for only one active link. An active link cannot belong to another FlexLinks pair.
- Neither of the links can be a port that belongs to an EtherChannel. However, you can configure two port channels (EtherChannel logical interfaces) as FlexLinks, and you can configure a port channel and a physical interface as FlexLinks, with either the port channel or the physical interface as the active link.
- A backup link does not have to be the same type (Fast Ethernet, Gigabit Ethernet, or port channel) as the active link. However, you should configure both FlexLinks with similar characteristics so that there are no loops or changes in behavior if the standby link begins to forward traffic.
- STP is disabled on FlexLinks ports. A FlexLinks port does not participate in STP, even if the VLANs present on the port are configured for STP. When STP is not enabled, be sure that there are no loops in the configured topology. Once the FlexLinks configurations are removed, STP is reenabled on the ports.

Follow these guidelines to configure VLAN load balancing on the FlexLinks feature:

- For FlexLinks VLAN load balancing, you must choose the preferred VLANs on the backup interface.
- You cannot configure a preemption mechanism and VLAN load balancing for the same FlexLinks pair.

Follow these guidelines to configure the MAC address-table move update feature:

- You can enable and configure this feature on the access switch to *send* the MAC address-table move updates.
- You can enable and configure this feature on the uplink switches to *receive* the MAC address-table move updates.

How to Configure the FlexLinks and MAC Address-Table Move Update

Configuring FlexLinks

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface, and enters interface configuration mode. The interface can be a physical Layer 2 interface or a port channel (logical interface). The port-channel range is 1 to 6.

	Command	Purpose
Step 3	**switchport backup interface** *interface-id*	Configures a physical Layer 2 interface (or port channel) as part of a FlexLinks pair with the interface. When one link is forwarding traffic, the other interface is in standby mode.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring a Preemption Scheme for FlexLinks

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface, and enter interface configuration mode. The interface can be a physical Layer 2 interface or a port channel (logical interface). The port-channel range is 1 to 6.
Step 3	**switchport backup interface** *interface-id*	Configures a physical Layer 2 interface (or port channel) as part of a FlexLinks pair with the interface. When one link is forwarding traffic, the other interface is in standby mode.
Step 4	**switchport backup interface** *interface-id* **preemption mode** [**forced** | **bandwidth** | **off**]	Configures a preemption mechanism and delay for a FlexLinks interface pair. You can configure the preemption as: • **forced**—The active interface always preempts the backup. • **bandwidth**—The interface with the higher bandwidth always acts as the active interface. • **off**—No preemption happens from active to backup.
Step 5	**switchport backup interface** *interface-id* **preemption delay** *delay-time*	Configures the time delay until a port preempts another port. **Note** Setting a delay time only works with forced and bandwidth modes.
Step 6	**end**	Returns to privileged EXEC mode.

Configuring VLAN Load Balancing on FlexLinks

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface, and enters interface configuration mode. The interface can be a physical Layer 2 interface or a port channel (logical interface). The port-channel range is 1 to 6.

	Command	Purpose
Step 3	**switchport backup interface** *interface-id* ***prefer vlan*** *vlan-range*	Configures a physical Layer 2 interface (or port channel) as part of a FlexLinks pair with the interface, and specifies the VLANs carried on the interface. The VLAN ID range is 1 to 4096.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring the MAC Address-Table Move Update Feature

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface, and enters interface configuration mode. The interface can be a physical Layer 2 interface or a port channel (logical interface). The port-channel range is 1 to 6.
Step 3	**switchport backup interface** *interface-id* or **switchport backup interface** *interface-id* **mmu primary vlan** *vlan-id*	Configures a physical Layer 2 interface (or port channel), as part of a FlexLinks pair with the interface. The MAC address-table move update VLAN is the lowest VLAN ID on the interface. Configures a physical Layer 2 interface (or port channel) and specifies the VLAN ID on the interface, which is used for sending the MAC address-table move update. When one link is forwarding traffic, the other interface is in standby mode.
Step 4	**end**	Returns to global configuration mode.
Step 5	**mac address-table move update transmit**	Enables the access switch to send MAC address-table move updates to other switches in the network if the primary link goes down and the switch starts forwarding traffic through the standby link.
Step 6	**end**	Returns to privileged EXEC mode.

Configuring the MAC Address-Table Move Update Messages

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mac address-table move update receive**	Enables the switch to get and process the MAC address-table move updates.
Step 3	**end**	Returns to privileged EXEC mode.
Step 4	**show mac address-table move update**	Verifies the configuration.
Step 5	**copy running-config startup config**	(Optional) Saves your entries in the switch startup configuration file.

Maintaining and Monitoring the FlexLinks and MAC Address-Table Move Update

Command	Purpose
show interfaces [*interface-id*] **switchport backup**	Displays the FlexLinks backup interface configured for an interface or all the configured FlexLinks and the state of each active and backup interface (up or standby mode). When VLAN load balancing is enabled, the output displays the preferred VLANs on active and backup interfaces.
show mac address-table move update	Verifies the configuration.

Configuration Examples for the FlexLinks and MAC Address-Table Move Update

Configuring FlexLinks Port: Examples

These are configuration examples for learning the other FlexLinks port as the mrouter port when FlexLinks is configured, with output for the **show interfaces switchport backup** command:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface GigabitEthernet1/1
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport mode trunk
Switch(config-if)# switchport backup interface GigabitEthernet1/2
Switch(config-if)# exit
Switch(config)# interface GigabitEthernet1/2
Switch(config-if)# switchport mode trunk
Switch(config-if)# end
Switch# show interfaces switchport backup detail
Switch Backup Interface Pairs:
Active Interface Backup Interface State
Preemption Mode : off
Multicast Fast Convergence : Off
Mac Address Move Update Vlan : auto
```

This output shows a querier for VLANs 1 and 401, with their queries reaching the switch through the specified port:

```
Switch# show ip igmp snooping querier
Vlan    IP Address      IGMP Version       Port
-----------------------------------------------------------
1       1.1.1.1         v2                 Gi0/1
401     41.41.41.1      v2                 Gi0/1
```

Here is output for the **show ip igmp snooping mrouter** command for VLANs 1 and 401:

```
Switch# show ip igmp snooping mrouter
Vlan    ports
```

```
----    -----
1       Gi1/1(dynamic), Gi1/2(dynamic)
401     Gi1/1(dynamic), Gi1/2(dynamic)
```

Similarly, both FlexLinks ports are part of learned groups. In this example, GigabitEthernet1/1 is a receiver/host in VLAN 1, which is interested in two multicast groups:

```
Switch# show ip igmp snooping groups
Vlan    Group       Type    Version     Port List
------------------------------------------------------------------------
1       228.1.5.1   igmp    v2          Gi1/1, Gi1/2, Fa2/1
1       228.1.5.2   igmp    v2          Gi1/1, Gi1/2, Fa2/1
```

When a host responds to the general query, the switch forwards this report on all the mrouter ports. In this example, when a host sends a report for the group 228.1.5.1, it is forwarded only on GigabitEthernet1/1, because the backup port GigabitEthernet1/2 is blocked. When the active link, GigabitEthernet1/1, goes down, the backup port, GigabitEthernet1/2, begins forwarding.

As soon as this port starts forwarding, the switch sends proxy reports for the groups 228.1.5.1 and 228.1.5.2 on behalf of the host. The upstream router learns the groups and starts forwarding multicast data. This is the default behavior of FlexLinks. This behavior changes when the user configures fast convergence using the **switchport backup interface GigabitEthernet**1/2 **multicast fast-convergence** command. This example shows how this feature is configured:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface GigabitEthernet1/1
Switch(config-if)# switchport backup interface GigabitEthernet1/2 multicast
fast-convergence
Switch(config-if)# exit
Switch# show interfaces switchport backup detail

Switch Backup Interface Pairs:
Active                Interface             Backup Interface State
------------------------------------------------------------------------
GigabitEthernet1/1  GigabitEthernet1/2  Active Up/Backup Standby
Preemption Mode : off
Multicast Fast Convergence : On
Mac Address Move Update Vlan : auto
```

This output shows a querier for VLAN 1 and 401 with their queries reaching the switch through the configured port:

```
Switch# show ip igmp snooping querier
Vlan       IP Address     IGMP Version   Port
-----------------------------------------------------------
1          1.1.1.1        v2             Gi1/1
401        41.41.41.1     v2             Gi1/1
```

This is output for the **show ip igmp snooping mrouter** command for VLAN 1 and 401:

```
Switch# show ip igmp snooping mrouter
Vlan       ports
----       -----
1          Gi1/1(dynamic), Gi1/2(dynamic)
401        Gi1/1(dynamic), Gi1/2(dynamic)
```

Similarly, both the FlexLinks ports are a part of the learned groups. In this example, the port is a receiver/host in VLAN 1, which is interested in two multicast groups:

```
Switch# show ip igmp snooping groups
Vlan    Group       Type    Version     Port List
------------------------------------------------------------------------
```

```
1       228.1.5.1  igmp    v2          Gi1/1, Gi1/2, Gi1/1
1       228.1.5.2  igmp    v2          Gi1/1, Gi1/2, Gi1/1
```

Whenever a host responds to the general query, the switch forwards this report on all the mrouter ports. When you turn on this feature through the command-line port, and when a report is forwarded by the switch on the configured GigabitEthernet1/1, it is also leaked to the backup port GigabitEthernet1/2. The upstream router learns the groups and starts forwarding multicast data, which is dropped at the ingress because the GigabitEthernet1/2 is blocked. When the active link, GigabitEthernet1/1 goes down, the backup port, GigabitEthernet1/2, begins forwarding. You do not need to send any proxy reports because the multicast data is already being forwarded by the upstream router. By leaking reports to the backup port, a redundant multicast path has been set up, and the time taken for the multicast traffic convergence is minimal.

Configuring a Backup Interface: Example

This example shows how to configure an interface with a backup interface and to verify the configuration:

```
Switch# configure terminal
Switch(conf)# interface gigabitethernet1/1
Switch(conf-if)# switchport backup interface gigabitethernet1/2
Switch(conf-if)# end

Switch# show interfaces switchport backup
Switch Backup Interface Pairs:

Active Interface        Backup Interface        State
------------------------------------------------------------------------
Vlans Preferred on Active Interface: 1-3,5-4096
        Vlans Preferred on Backup Interface: 4
```

Configuring a Preemption Scheme: Example

This example shows how to configure the preemption mode as *forced* for a backup interface pair and to verify the configuration:

```
Switch# configure terminal
Switch(conf)# interface gigabitethernet1/1
Switch(conf-if)#switchport backup interface gigabitethernet1/2 preemption mode forced
Switch(conf-if)#switchport backup interface gigabitethernet1/2 preemption delay 50
Switch(conf-if)# end

Switch# show interfaces switchport backup detail
Active Interface Backup Interface State
------------------------------------------------------------------------
GigabitEthernet1/1 GigabitEthernet1/2 Active Up/Backup Standby
Interface Pair : Gi1/1, Gi1/2
Preemption Mode : forced
Preemption Delay : 50 seconds
Bandwidth : 100000 Kbit (Gi1/1), 100000 Kbit (Gi1/2)
Mac Address Move Update Vlan : auto
```

Configuring VLAN Load Balancing on FlexLinks: Examples

In the following example, VLANs 1 to 50, 60, and 100 to 120 are configured on the switch:

```
Switch(config)# interface gigabitEthernet 1/2
Switch(config-if)# switchport backup interface gigabitEthernet 1/2 prefer vlan 60,100-120
```

When both interfaces are up, GigabitEthernet1/1 forwards traffic for VLANs 60 and 100 to 120, and GigabitEthernet1/2 forwards traffic for VLANs 1 to 50.

```
Switch# show interfaces switchport backup
Switch Backup Interface Pairs:

Active Interface        Backup Interface        State
------------------------------------------------------------------------
GigabitEthernet1/1    GigabitEthernet1/2    Active Up/Backup Standby
Vlans Preferred on Active Interface: 1-50
Vlans Preferred on Backup Interface: 60, 100-120
```

When a FlexLinks interface goes down (LINK_DOWN), VLANs preferred on this interface are moved to the peer interface of the FlexLinks pair. In this example, if interface Gigabit Ethernet1/1 goes down, Gigabit Ethernet1/2 carries all VLANs of the FlexLinks pair.

```
Switch# show interfaces switchport backup
Switch Backup Interface Pairs:

Active Interface        Backup Interface        State
------------------------------------------------------------------------
GigabitEthernet1/1    GigabitEthernet1/2    Active Down/Backup Up

Vlans Preferred on Active Interface: 1-50
Vlans Preferred on Backup Interface: 60, 100-120
```

When a FlexLinks interface comes up, VLANs preferred on this interface are blocked on the peer interface and moved to the forwarding state on the interface that has just come up. In this example, if interface Gigabit Ethernet1/1 comes up, VLANs preferred on this interface are blocked on the peer interface Gigabit Ethernet1/2 and forwarded on Gigabit Ethernet1/1.

```
Switch# show interfaces switchport backup
Switch Backup Interface Pairs:

Active Interface        Backup Interface        State
------------------------------------------------------------------------
GigabitEthernet1/1    GigabitEthernet1/2    Active Down/Backup Up

Vlans Preferred on Active Interface: 1-50
Vlans Preferred on Backup Interface: 60, 100-120

Switch# show interfaces switchport backup detail
Switch Backup Interface Pairs:

Active Interface        Backup Interface        State
------------------------------------------------------------------------
FastEthernet1/3       FastEthernet1/4       Active Down/Backup Up

Vlans Preferred on Active Interface: 1-2,5-4096
Vlans Preferred on Backup Interface: 3-4
Preemption Mode  : off
Bandwidth : 10000 Kbit (Fa1/3), 100000 Kbit (Fa1/4)
Mac Address Move Update Vlan : auto
```

Configuring MAC Address-Table Move Update: Example

This example shows how to configure an access switch to send MAC address-table move update messages:

```
Switch(conf)# interface gigabitethernet1/1
Switch(conf-if)# switchport backup interface gigabitethernet1/2 mmu primary vlan 2
Switch(conf-if)# exit
Switch(conf)# mac address-table move update transmit
Switch(conf)# end
```

This example shows how to verify the configuration:

```
Switch# show mac-address-table move update
Switch-ID : 010b.4630.1780
Dst mac-address : 0180.c200.0010
Vlans/Macs supported : 1023/8320
Default/Current settings: Rcv Off/On, Xmt Off/On
Max packets per min : Rcv 40, Xmt 60
Rcv packet count : 5
Rcv conforming packet count : 5
Rcv invalid packet count : 0
Rcv packet count this min : 0
Rcv threshold exceed count : 0
Rcv last sequence# this min : 0
Rcv last interface : Po2
Rcv last src-mac-address : 000b.462d.c502
Rcv last switch-ID : 0403.fd6a.8700
Xmt packet count : 0
Xmt packet count this min : 0
Xmt threshold exceed count : 0
Xmt pak buf unavail cnt : 0
Xmt last interface : None
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

CHAPTER 25

Configuring DHCP

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Information About Configuring DHCP

This chapter describes how to configure Dynamic Host Configuration Protocol (DHCP) snooping and option-82 data insertion, and the DHCP server port-based address allocation features on the switch. It also describes how to configure the IP source guard feature.

DHCP Snooping

DHCP is widely used in LAN environments to dynamically assign host IP addresses from a centralized server, which significantly reduces the overhead of administration of IP addresses. DHCP also helps conserve the limited IP address space because IP addresses no longer need to be permanently assigned to hosts; only those hosts that are connected to the network consume IP addresses.

DHCP Server

The DHCP server assigns IP addresses from specified address pools on a switch or router to DHCP clients and manages them. If the DHCP server cannot give the DHCP client the requested configuration parameters from its database, it forwards the request to one or more secondary DHCP servers defined by the network administrator.

DHCP Relay Agent

A DHCP relay agent is a Layer 3 device that forwards DHCP packets between clients and servers. Relay agents forward requests and replies between clients and servers when they are not on the same physical subnet. Relay agent forwarding is different from the normal Layer 2 forwarding, in which IP datagrams are switched transparently between networks. Relay agents receive DHCP messages and generate new DHCP messages to send on output interfaces.

DHCP Snooping

DHCP snooping is a DHCP security feature that provides network security by filtering untrusted DHCP messages and by building and maintaining a DHCP snooping binding database, also referred to as a DHCP snooping binding table.

DHCP snooping acts like a firewall between untrusted hosts and DHCP servers. You use DHCP snooping to differentiate between untrusted interfaces connected to the end user and trusted interfaces connected to the DHCP server or another switch.

Note For DHCP snooping to function properly, all DHCP servers must be connected to the switch through trusted interfaces.

An untrusted DHCP message is a message that is received from outside the network or firewall. When you use DHCP snooping in a service-provider environment, an untrusted message is sent from a device that is not in the service-provider network, such as a customer's switch. Messages from unknown devices are untrusted because they can be sources of traffic attacks.

The DHCP snooping binding database has the MAC address, the IP address, the lease time, the binding type, the VLAN number, and the interface information that corresponds to the local untrusted interfaces of a switch. It does not have information regarding hosts interconnected with a trusted interface.

In a service-provider network, a trusted interface is connected to a port on a device in the same network. An untrusted interface is connected to an untrusted interface in the network or to an interface on a device that is not in the network.

When a switch receives a packet on an untrusted interface and the interface belongs to a VLAN in which DHCP snooping is enabled, the switch compares the source MAC address and the DHCP client hardware address. If the addresses match (the default), the switch forwards the packet. If the addresses do not match, the switch drops the packet.

The switch drops a DHCP packet when one of these situations occurs:

- A packet from a DHCP server, such as a DHCPOFFER, DHCPACK, DHCPNAK, or DHCPLEASEQUERY packet, is received from outside the network or firewall.
- A packet is received on an untrusted interface, and the source MAC address and the DHCP client hardware address do not match.
- The switch receives a DHCPRELEASE or DHCPDECLINE broadcast message that has a MAC address in the DHCP snooping binding database, but the interface information in the binding database does not match the interface on which the message was received.
- A DHCP relay agent forwards a DHCP packet that includes a relay-agent IP address that is not 0.0.0.0, or the relay agent forwards a packet that includes option-82 information to an untrusted port.

If the switch is an aggregation switch supporting DHCP snooping and is connected to an edge switch that is inserting DHCP option-82 information, the switch drops packets with option-82 information when packets are received on an untrusted interface. If DHCP snooping is enabled and packets are received on a trusted port, the aggregation switch does not learn the DHCP snooping bindings for connected devices and cannot build a complete DHCP snooping binding database.

When an aggregation switch can be connected to an edge switch through an untrusted interface and you enter the **ip dhcp snooping information option allow-untrusted** global configuration command, the aggregation switch accepts packets with option-82 information from the edge switch. The aggregation switch learns the bindings for hosts connected through an untrusted switch interface. The DHCP security features, such as dynamic ARP inspection or IP source guard, can still be enabled on the aggregation switch while the switch receives packets with option-82 information on untrusted input interfaces to which hosts are connected. The port on the edge switch that connects to the aggregation switch must be configured as a trusted interface.

Option-82 Data Insertion

In residential, metropolitan Ethernet-access environments, DHCP can centrally manage the IP address assignments for a large number of subscribers. When the DHCP option-82 feature is enabled on the switch, a subscriber device is identified by the switch port through which it connects to the network (in addition to its MAC address). Multiple hosts on the subscriber LAN can be connected to the same port on the access switch and are uniquely identified.

Note The DHCP option-82 feature is supported only when DHCP snooping is globally enabled and on the VLANs to which subscriber devices using this feature are assigned.

Figure 25-1 is an example of a metropolitan Ethernet network in which a centralized DHCP server assigns IP addresses to subscribers connected to the switch at the access layer. Because the DHCP clients and their associated DHCP server do not reside on the same IP network or subnet, a DHCP relay agent (the Catalyst switch) is configured with a helper address to enable broadcast forwarding and to transfer DHCP messages between the clients and the server.

Figure 25-1 DHCP Relay Agent in a Metropolitan Ethernet Network

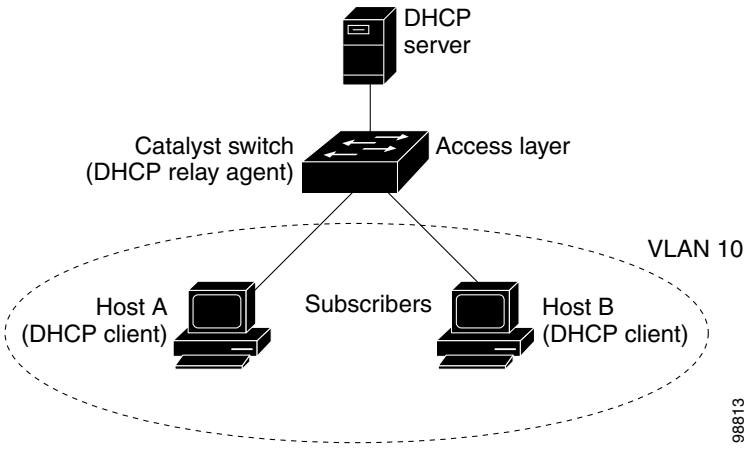

When you enable the DHCP snooping information option-82 on the switch, this sequence of events occurs:

- The host (DHCP client) generates a DHCP request and broadcasts it on the network.
- When the switch receives the DHCP request, it adds the option-82 information in the packet. By default, the remote-ID suboption is the switch MAC address, and the circuit-ID suboption is the port identifier, **vlan-mod-port**, from which the packet is received.
- If the IP address of the relay agent is configured, the switch adds this IP address in the DHCP packet.
- The switch forwards the DHCP request that includes the option-82 field to the DHCP server.
- The DHCP server receives the packet. If the server is option-82-capable, it can use the remote ID, the circuit ID, or both to assign IP addresses and implement policies, such as restricting the number of IP addresses that can be assigned to a single remote ID or circuit ID. The DHCP server then repeats the option-82 field in the DHCP reply.
- The DHCP server unicasts the reply to the switch if the request was relayed to the server by the switch. The switch verifies that it originally inserted the option-82 data by inspecting the remote ID and possibly the circuit ID fields. The switch removes the option-82 field and forwards the packet to the switch port that connects to the DHCP client that sent the DHCP request.

In the default suboption configuration, when the described sequence of events occurs, the values in these fields in Figure 25-2 do not change:

- Circuit-ID suboption fields
  - Suboption type
  - Length of the suboption type
  - Circuit-ID type
  - Length of the circuit-ID type
- Remote-ID suboption fields
  - Suboption type
  - Length of the suboption type
  - Remote-ID type
  - Length of the remote-ID type

In the port field of the circuit-ID suboption, the port numbers start at 3. For example, on a switch with eight 10/100 ports and small form-factor pluggable (SFP) module slots, port 3 is the Fast Ethernet 1/1 port, port 4 is the Fast Ethernet 1/2 port, and so forth. Port 11 is the SFP module slot 1/1, and so forth.

Figure 25-2 shows the packet formats for the remote-ID suboption and the circuit-ID suboption when the default suboption configuration is used. The switch uses the packet formats when you globally enable DHCP snooping and enter the **ip dhcp snooping information option** global configuration command.

Figure 25-2 Suboption Packet Formats

Circuit ID Suboption Frame Format

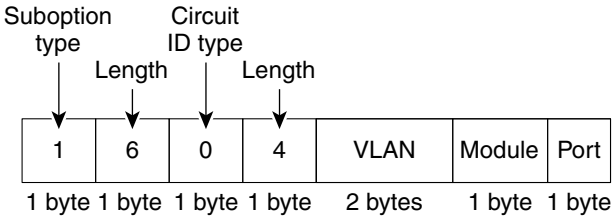

Remote ID Suboption Frame Format

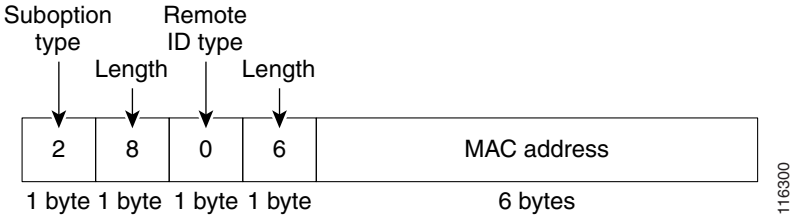

Figure 25-3 shows the packet formats for user-configured remote-ID and circuit-ID suboptions The switch uses these packet formats when DHCP snooping is globally enabled and when the **ip dhcp snooping information option format remote-id** global configuration command **and the ip dhcp snooping vlan information option format-type circuit-id string** interface configuration command are entered.

The values for these fields in the packets change from the default values when you configure the remote-ID and circuit-ID suboptions:

- Circuit-ID suboption fields
  - The circuit-ID type is 1.
  - The length values are variable, depending on the length of the string that you configure.
- Remote-ID suboption fields
  - The remote-ID type is 1.
  - The length values are variable, depending on the length of the string that you configure.

Figure 25-3 User-Configured Suboption Packet Formats

Circuit ID Suboption Frame Format (for user-configured string):

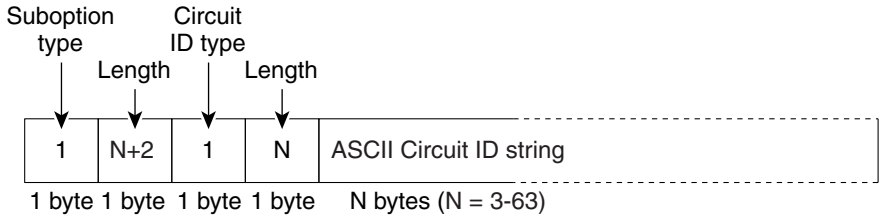

Remote ID Suboption Frame Format (for user-configured string):

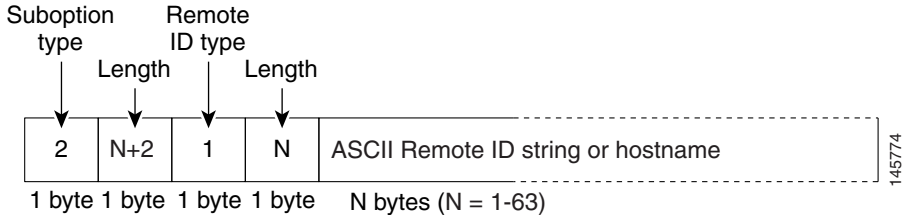

Cisco IOS DHCP Server Database

During the DHCP-based autoconfiguration process, the designated DHCP server uses the Cisco IOS DHCP server database. It has IP addresses, *address bindings*, and configuration parameters, such as the boot file.

An address binding is a mapping between an IP address and a MAC address of a host in the Cisco IOS DHCP server database. You can manually assign the client IP address, or the DHCP server can allocate an IP address from a DHCP address pool.

DHCP Snooping Binding Database

When DHCP snooping is enabled, the switch uses the DHCP snooping binding database to store information about untrusted interfaces. The database can have up to 8192 bindings.

Each database entry (*binding*) has an IP address, an associated MAC address, the lease time (in hexadecimal format), the interface to which the binding applies, and the VLAN to which the interface belongs. The database agent stores the bindings in a file at a configured location. At the end of each entry is a checksum that accounts for all the bytes from the start of the file through all the bytes associated with the entry. Each entry is 72 bytes, followed by a space and then the checksum value.

To keep the bindings when the switch reloads, you must use the DHCP snooping database agent. If the agent is disabled, dynamic ARP inspection or IP source guard is enabled, and the DHCP snooping binding database has dynamic bindings, the switch loses its connectivity. If the agent is disabled and only DHCP snooping is enabled, the switch does not lose its connectivity, but DHCP snooping might not prevent DHCP spoofing attacks.

When reloading, the switch reads the binding file to build the DHCP snooping binding database. The switch updates the file when the database changes.

When a switch learns of new bindings or when it loses bindings, the switch immediately updates the entries in the database. The switch also updates the entries in the binding file. The frequency at which the file is updated is based on a configurable delay, and the updates are batched. If the file is not updated in a specified time (set by the write-delay and abort-timeout values), the update stops.

This is the format of the file with bindings:

```
<initial-checksum>
TYPE DHCP-SNOOPING
VERSION 1
BEGIN
<entry-1> <checksum-1>
<entry-2> <checksum-1-2>
...
...
<entry-n> <checksum-1-2-..-n>
END
```

Each entry in the file is tagged with a checksum value that the switch uses to verify the entries when it reads the file. The *initial-checksum* entry on the first line distinguishes entries associated with the latest file update from entries associated with a previous file update.

This is an example of a binding file:

```
2bb4c2a1
TYPE DHCP-SNOOPING
VERSION 1
BEGIN
192.1.168.1 3 0003.47d8.c91f 2BB6488E interface-id 21ae5fbb
192.1.168.3 3 0003.44d6.c52f 2BB648EB interface-id 1bdb223f
192.1.168.2 3 0003.47d9.c8f1 2BB648AB interface-id 584a38f0
END
```

When the switch starts and the calculated checksum value equals the stored checksum value, the switch reads entries from the binding file and adds the bindings to its DHCP snooping binding database. The switch ignores an entry when one of these situations occurs:

- The switch reads the entry and the calculated checksum value does not equal the stored checksum value. The entry and the ones following it are ignored.
- An entry has an expired lease time (the switch might not remove a binding entry when the lease time expires).
- The interface in the entry no longer exists on the system.
- The interface is a routed interface or a DHCP snooping-trusted interface.

Default DHCP Snooping Settings

Table 25-1 Default DHCP Snooping Settings

Feature	Default Setting
DHCP server	Enabled in Cisco IOS software, requires configuration[1]
DHCP relay agent	Enabled[2]
DHCP packet forwarding address	None configured
Checking the relay agent information	Enabled (invalid messages are dropped)[2]
DHCP relay agent forwarding policy	Replace the existing relay agent information[2]

Table 25-1 Default DHCP Snooping Settings (continued)

Feature	Default Setting
DHCP snooping enabled globally	Disabled
DHCP snooping information option	Enabled
DHCP snooping option to accept packets on untrusted input interfaces[3]	Disabled
DHCP snooping limit rate	None configured
DHCP snooping trust	Untrusted
DHCP snooping VLAN	Disabled
DHCP snooping MAC address verification	Enabled
Cisco IOS DHCP server binding database	Enabled in Cisco IOS software, requires configuration. **Note** The switch gets network addresses and configuration parameters only from a device configured as a DHCP server.
DHCP snooping binding database agent	Enabled in Cisco IOS software, requires configuration. This feature is operational only when a destination is configured.

1. The switch responds to DHCP requests only if it is configured as a DHCP server.
2. The switch relays DHCP packets only if the IP address of the DHCP server is configured on the SVI of the DHCP client.
3. Use this feature when the switch is an aggregation switch that receives packets with option-82 information from an edge switch.

DHCP Snooping Configuration Guidelines

- You must globally enable DHCP snooping on the switch.
- DHCP snooping is not active until DHCP snooping is enabled on a VLAN.
- Before globally enabling DHCP snooping on the switch, make sure that the devices acting as the DHCP server and the DHCP relay agent are configured and enabled.
- Before configuring the DHCP snooping information option on your switch, be sure to configure the device that is acting as the DHCP server. For example, you must specify the IP addresses that the DHCP server can assign or exclude, or you must configure DHCP options for these devices.
- When configuring a large number of circuit IDs on a switch, consider the impact of lengthy character serstrings on the NVRAM or the flash memory. If the circuit-ID configurations, combined with other data, exceed the capacity of the NVRAM or the flash memory, an error message appears.
- Before configuring the DHCP relay agent on your switch, make sure to configure the device that is acting as the DHCP server. For example, you must specify the IP addresses that the DHCP server can assign or exclude, configure DHCP options for devices, or set up the DHCP database agent.
- If the DHCP relay agent is enabled but DHCP snooping is disabled, the DHCP option-82 data insertion feature is not supported.
- If a switch port is connected to a DHCP server, configure a port as trusted by entering the **ip dhcp snooping trust** interface configuration command.
- If a switch port is connected to a DHCP client, configure a port as untrusted by entering the **no ip dhcp snooping trust** interface configuration command.
- Do not enter the **ip dhcp snooping information option allow-untrusted** command on an aggregation switch to which an untrusted device is connected. If you enter this command, an untrusted device might spoof the option-82 information.

- You can display DHCP snooping statistics by entering the **show ip dhcp snooping statistics** user EXEC command, and you can clear the snooping statistics counters by entering the **clear ip dhcp snooping statistics** privileged EXEC command.

Note Do not enable DHCP snooping on RSPAN VLANs. If DHCP snooping is enabled on RSPAN VLANs, DHCP packets might not reach the RSPAN destination port.

DHCP Snooping Binding Database Guidelines

- Because both NVRAM and the flash memory have limited storage capacity, we recommend that you store the binding file on a TFTP server.
- For network-based URLs (such as TFTP and FTP), you must create an empty file at the configured URL before the switch can write bindings to the binding file at that URL. See the documentation for your TFTP server to determine whether you must first create an empty file on the server; some TFTP servers cannot be configured this way.
- To ensure that the lease time in the database is accurate, we recommend that you enable and configure NTP. For more information, see the "Configuring Time and Date Manually" section on page 7-9.
- If NTP is configured, the switch writes binding changes to the binding file only when the switch system clock is synchronized with NTP.

Packet Forwarding Address

If the DHCP server and the DHCP clients are on different networks or subnets, you must configure the switch with the **ip helper-address** *address* interface configuration command. The general rule is to configure the command on the Layer 3 interface closest to the client. The address used in the **ip helper-address** command can be a specific DHCP server IP address, or it can be the network address if other DHCP servers are on the destination network segment. Using the network address enables any DHCP server to respond to requests.

DHCP Server Port-Based Address Allocation

DHCP server port-based address allocation is a feature that enables DHCP to maintain the same IP address on an Ethernet switch port regardless of the attached device client identifier or client hardware address.

When Ethernet switches are deployed in the network, they offer connectivity to the directly connected devices. In some environments, such as on a factory floor, if a device fails, the replacement device must be working immediately in the existing network. With the current DHCP implementation, there is no guarantee that DHCP would offer the same IP address to the replacement device. Control, monitoring, and other software expect a stable IP address associated with each device. If a device is replaced, the address assignment should remain stable even though the DHCP client has changed.

When configured, the DHCP server port-based address allocation feature ensures that the same IP address is always offered to the same connected port even as the client identifier or client hardware address changes in the DHCP messages received on that port. The DHCP protocol recognizes DHCP clients by the client identifier option in the DHCP packet. Clients that do not include the client identifier

option are identified by the client hardware address. When you configure this feature, the port name of the interface overrides the client identifier or hardware address and the actual point of connection, the switch port, becomes the client identifier.

In all cases, by connecting the Ethernet cable to the same port, the same IP address is allocated through DHCP to the attached device.

The DHCP server port-based address allocation feature is only supported on a Cisco IOS DHCP server and not a third-party server.

By default, DHCP server port-based address allocation is disabled.

How to Configure DHCP

Configuring the DHCP Relay Agent

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**service dhcp**	Enables the DHCP server and relay agent on your switch. By default, this feature is enabled.
Step 3	**end**	Returns to privileged EXEC mode.

Specifying the Packet Forwarding Address

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface vlan** *vlan-id*	Creates a switch virtual interface by entering a VLAN ID, and enters interface configuration mode.
Step 3	**ip address** *ip-address subnet-mask*	Configures the interface with an IP address and an IP subnet.
Step 4	**ip helper-address** *address*	Specifies the DHCP packet forwarding address. The helper address can be a specific DHCP server address, or it can be the network address if other DHCP servers are on the destination network segment. Using the network address enables other servers to respond to DHCP requests. If you have multiple servers, you can configure one helper address for each server.
Step 5	**exit**	Returns to global configuration mode.

	Command	Purpose
Step 6	**interface range** *port-range*	Configures multiple physical ports that are connected to the DHCP clients, and enters interface range configuration mode.
	or	or
	interface *interface-id*	Configures a single physical port that is connected to the DHCP client, and enters interface configuration mode.
Step 7	**switchport mode access**	Defines the VLAN membership mode for the port.
Step 8	**switchport access vlan** *vlan-id*	Assigns the ports to the same VLAN as configured in Step 2.
Step 9	**end**	Returns to privileged EXEC mode.

Enabling DHCP Snooping and Option 82

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip dhcp snooping**	Enables DHCP snooping globally.
Step 3	**ip dhcp snooping vlan** *vlan-range*	Enables DHCP snooping on a VLAN or range of VLANs. The range is 1 to 4096. You can enter a single VLAN ID identified by VLAN ID number, a series of VLAN IDs separated by commas, a range of VLAN IDs separated by hyphens, or a range of VLAN IDs separated by entering the starting and ending VLAN IDs separated by a space.
Step 4	**ip dhcp snooping information option**	Enables the switch to insert and to remove DHCP relay information (option-82 field) in forwarded DHCP request messages to the DHCP server. This is the default setting.
Step 5	**ip dhcp snooping information option format remote-id [string** *ASCII-string* | *hostname*]	(Optional) Configures the remote-ID suboption. You can configure the remote ID as • String of up to 63 ASCII characters (no spaces) • Hostname for the switch **Note** If the hostname is longer than 63 characters, it is truncated to 63 characters in the remote-ID configuration. The default remote ID is the switch MAC address.
Step 6	**ip dhcp snooping information option allow-untrusted**	(Optional) If the switch is an aggregation switch connected to an edge switch, enable the switch to accept incoming DHCP snooping packets with option-82 information from the edge switch. The default setting is disabled. **Note** Enter this command only on aggregation switches that are connected to trusted devices.

	Command	Purpose
Step 7	**interface** *interface-id*	Specifies the interface to be configured, and enters interface configuration mode.
Step 8	**ip dhcp snooping vlan** *vlan* **information option format-type circuit-id** [**override**] **string** *ASCII-string*	(Optional) Configures the circuit-ID suboption for the specified interface. Specifies the VLAN and port identifier, using a VLAN ID in the range of 1 to 4096. The default circuit ID is the port identifier in the format **vlan-mod-port.** You can configure the circuit ID to be a string of 3 to 63 ASCII characters (no spaces). (Optional) Use the **override** keyword when you do not want the circuit-ID suboption inserted in TLV format to define subscriber information.
Step 9	**ip dhcp snooping trust**	(Optional) Configures the interface as trusted or as untrusted. Use the **no** keyword to configure an interface to receive messages from an untrusted client. The default setting is untrusted.
Step 10	**ip dhcp snooping limit rate** *rate*	(Optional) Configures the number of DHCP packets per second that an interface can receive. The range is 1 to 2048. By default, no rate limit is configured. **Note** We recommend an untrusted rate limit of not more than 100 packets per second. If you configure rate limiting for trusted interfaces, you might need to increase the rate limit if the port is a trunk port assigned to more than one VLAN with DHCP snooping.
Step 11	**exit**	Returns to global configuration mode.
Step 12	**ip dhcp snooping verify mac-address**	(Optional) Configures the switch to verify that the source MAC address in a DHCP packet received on untrusted ports matches the client hardware address in the packet. The default is to verify that the source MAC address matches the client hardware address in the packet.
Step 13	**end**	Returns to privileged EXEC mode.

Enabling the DHCP Snooping Binding Database Agent

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip dhcp snooping database** {**flash:**/*filename* | **ftp://***user***:***password***@***host***/***filename* | **http://**[[*username***:***password*]@]{*hostname* | *host-ip*}[/*directory*] /*image-name*.**tar** | **rcp://***user***@***host***/***filename*}| **tftp://***host***/***filename*	Specifies the URL for the database agent or the binding file by using one of these forms: • **flash:/***filename* • **ftp://***user***:***password***@***host***/***filename* • **http://**[[*username***:***password*]@]{*hostname* | *host-ip*}[/*directory*] /*image-name*.**tar** • **rcp://***user***@***host***/***filename* • **tftp://***host***/***filename*

	Command	Purpose
Step 3	**ip dhcp snooping database timeout** *seconds*	Specifies (in seconds) how long to wait for the database transfer process to finish before stopping the process. The default is 300 seconds. The range is 0 to 86400. Use 0 to define an infinite duration, which means to continue trying the transfer indefinitely.
Step 4	**ip dhcp snooping database write-delay** *seconds*	Specifies the duration for which the transfer should be delayed after the binding database changes. The range is from 15 to 86400 seconds. The default is 300 seconds (5 minutes).
Step 5	**end**	Returns to privileged EXEC mode.
Step 6	**ip dhcp snooping binding** *mac-address* **vlan** *vlan-id ip-address* **interface** *interface-id* **expiry** *seconds*	(Optional) Adds binding entries to the DHCP snooping binding database. The *vlan-id* range is from 1 to 4904. The *seconds* range is from 1 to 4294967295. Enter this command for each entry that you add. **Note** Use this command when you are testing or debugging the switch.

Enabling DHCP Server Port-Based Address Allocation

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip dhcp use subscriber-id client-id**	Configures the DHCP server to globally use the subscriber identifier as the client identifier on all incoming DHCP messages.
Step 3	**ip dhcp subscriber-id interface-name**	Automatically generates a subscriber identifier based on the short name of the interface. A subscriber identifier configured on a specific interface takes precedence over this command.
Step 4	**interface** *interface-id*	Specifies the interface to be configured, and enters interface configuration mode.
Step 5	**ip dhcp server use subscriber-id client-id**	Configures the DHCP server to use the subscriber identifier as the client identifier on all incoming DHCP messages on the interface.
Step 6	**end**	Returns to privileged EXEC mode.

Preassigning an IP Address

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip dhcp pool** *poolname*	Enters DHCP pool configuration mode, and defines the name for the DHCP pool. The pool name can be a symbolic string (such as Engineering) or an integer (such as 0).

	Command	Purpose
Step 3	**network** *network-number* [*mask* | */prefix-length*]	Specifies the subnet network number and mask of the DHCP address pool.
Step 4	**address** *ip-address* **client-id** *string* [**ascii**]	Reserves an IP address for a DHCP client identified by the interface name. *string*—Can be an ASCII value or a hexadecimal value.
Step 5	**reserved-only**	(Optional) Uses only reserved addresses in the DHCP address pool. The default is to not restrict pool addresses.
Step 6	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining DHCP

Command	Purpose
show interface *interface id*	Displays the status and configuration of a specific interface.
show ip dhcp pool	Displays the DHCP address pools.
show ip dhcp binding	Displays address bindings on the Cisco IOS DHCP server.
ip dhcp snooping database timeout *seconds*	Specifies (in seconds) how long to wait for the database transfer process to finish before stopping.
ip dhcp snooping database write-delay *seconds*	Specifies (in seconds) the duration for which the transfer should be delayed after the binding database changes.
clear ip dhcp snooping database statistics	Clears the DHCP snooping binding database agent statistics.
renew ip dhcp snooping database	Renews the DHCP snooping binding database.
show ip dhcp snooping database [**detail**]	Displays the status and statistics of the DHCP snooping binding database agent.
show ip dhcp snooping	Displays the DHCP snooping configuration for a switch
show ip dhcp snooping binding	Displays only the dynamically configured bindings in the DHCP snooping binding database, also referred to as a binding table.
show ip dhcp snooping database	Displays the DHCP snooping binding database status and statistics.
show ip dhcp pool	Verifies DHCP pool configuration.
copy running-config startup-config	Saves your entries in the configuration file.

Configuration Examples for Configuring DHCP

Enabling DHCP Server Port-Based Address Allocation: Examples

In this example, a subscriber identifier is automatically generated, and the DHCP server ignores any client identifier fields in the DHCP messages and uses the subscriber identifier instead. The subscriber identifier is based on the short name of the interface and the client preassigned IP address 10.1.1.7.

```
switch# show running config
Building configuration...
Current configuration : 4899 bytes
!
version 12.2
!
hostname switch
!
no aaa new-model
clock timezone EST 0
ip subnet-zero
ip dhcp relay information policy removal pad
no ip dhcp use vrf connected
ip dhcp use subscriber-id client-id
ip dhcp subscriber-id interface-name
ip dhcp excluded-address 10.1.1.1 10.1.1.3
!
ip dhcp pool dhcppool
 network 10.1.1.0 255.255.255.0
 address 10.1.1.7 client-id “Et1/0” ascii
<output truncated>
```

This example shows that the preassigned address was correctly reserved in the DHCP pool:

```
switch# show ip dhcp pool dhcppool
Pool dhcp pool:
 Utilization mark (high/low) : 100 / 0
 Subnet size (first/next) : 0 / 0
 Total addresses : 254
 Leased addresses : 0
 Excluded addresses : 4
 Pending event : none
 1 subnet is currently in the pool:
 Current index   IP address range         Leased/Excluded/Total
 10.1.1.1        10.1.1.1 - 10.1.1.254     0     / 4 / 254
 1 reserved address is currently in the pool
 Address         Client
 10.1.1.7 Et1/0
```

Enabling DHCP Snooping: Example

This example shows how to enable DHCP snooping globally and on VLAN 10 and to configure a rate limit of 100 packets per second on a port:

```
Switch(config)# ip dhcp snooping
Switch(config)# ip dhcp snooping vlan 10
Switch(config)# ip dhcp snooping information option
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# ip dhcp snooping limit rate 100
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
Cisco IOS DHCP Commands	*Cisco IOS IP Command Reference, Volume 1 of 3: Addressing and Services*
Cisco IOS DHCP Configuration Cisco IOS DHCP server port-based address allocation	"IP Addressing and Services" chapter of the *Cisco IOS IP Configuration Guide*
Cisco IOS DHCP Configuration Task List	"Configuring DHCP" chapter of the *Cisco IOS IP Configuration Guide*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

CHAPTER 26

Configuring Dynamic ARP Inspection

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for Dynamic ARP Inspection

- Dynamic Address Resolution Protocol (ARP) inspection depends on the entries in the DHCP snooping binding database to verify IP-to-MAC address bindings in incoming ARP requests and ARP responses. Make sure to enable DHCP snooping to permit ARP packets that have dynamically assigned IP addresses.

Restrictions for Dynamic ARP Inspection

- To use this feature, the switch must be running the LAN Base image.

Information About Dynamic ARP Inspection

Dynamic ARP Inspection

Dynamic ARP inspection (DAI) helps prevent malicious attacks on the switch by not relaying invalid ARP requests and responses to other ports in the same VLAN.

ARP provides IP communication within a Layer 2 broadcast domain by mapping an IP address to a MAC address. For example, Host B wants to send information to Host A but does not have the MAC address of Host A in its ARP cache. Host B generates a broadcast message for all hosts within the broadcast domain to obtain the MAC address associated with the IP address of Host A. All hosts within the broadcast domain receive the ARP request, and Host A responds with its MAC address. However,

because ARP allows a gratuitous reply from a host even if an ARP request was not received, an ARP spoofing attack and the poisoning of ARP caches can occur. After the attack, all traffic from the device under attack flows through the attacker's computer and then to the router, switch, or host.

A malicious user can attack hosts, switches, and routers connected to your Layer 2 network by poisoning the ARP caches of systems connected to the subnet and by intercepting traffic intended for other hosts on the subnet. Figure 26-1 shows an example of ARP cache poisoning.

Figure 26-1 ARP Cache Poisoning

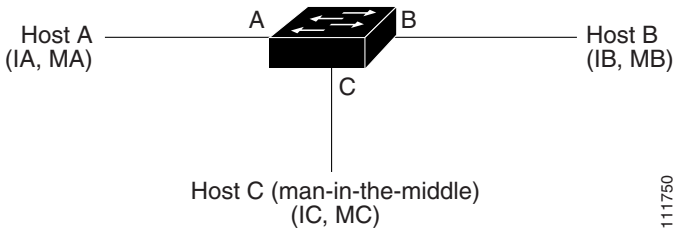

Hosts A, B, and C are connected to the switch on interfaces A, B and C, all of which are on the same subnet. Their IP and MAC addresses are shown in parentheses; for example, Host A uses IP address IA and MAC address MA. When Host A needs to communicate to Host B at the IP layer, it broadcasts an ARP request for the MAC address associated with IP address IB. When the switch and Host B receive the ARP request, they populate their ARP caches with an ARP binding for a host with the IP address IA and a MAC address MA; for example, IP address IA is bound to MAC address MA. When Host B responds, the switch and Host A populate their ARP caches with a binding for a host with the IP address IB and the MAC address MB.

Host C can poison the ARP caches of the switch, Host A, and Host B by broadcasting forged ARP responses with bindings for a host with an IP address of IA (or IB) and a MAC address of MC. Hosts with poisoned ARP caches use the MAC address MC as the destination MAC address for traffic intended for IA or IB. This means that Host C intercepts that traffic. Because Host C knows the true MAC addresses associated with IA and IB, it can forward the intercepted traffic to those hosts by using the correct MAC address as the destination. Host C has inserted itself into the traffic stream from Host A to Host B, the classic *man-in-the middle* attack.

DAI is a security feature that validates ARP packets in a network. It intercepts, logs, and discards ARP packets with invalid IP-to-MAC address bindings. This capability protects the network from certain man-in-the-middle attacks.

DAI ensures that only valid ARP requests and responses are relayed. The switch performs these activities:

- Intercepts all ARP requests and responses on untrusted ports
- Verifies that each of these intercepted packets has a valid IP-to-MAC address binding before updating the local ARP cache or before forwarding the packet to the appropriate destination
- Drops invalid ARP packets

DAI determines the validity of an ARP packet based on valid IP-to-MAC address bindings stored in a trusted database, the DHCP snooping binding database. This database is built by DHCP snooping if DHCP snooping is enabled on the VLANs and on the switch. If the ARP packet is received on a trusted interface, the switch forwards the packet without any checks. On untrusted interfaces, the switch forwards the packet only if it is valid.

Interface Trust States and Network Security

DAI associates a trust state with each interface on the switch. Packets arriving on trusted interfaces bypass all DAI validation checks, and those arriving on untrusted interfaces undergo the DAI validation process.

In a typical network configuration, you configure all switch ports connected to host ports as untrusted and configure all switch ports connected to switches as trusted. With this configuration, all ARP packets entering the network from a given switch bypass the security check. No other validation is needed at any other place in the VLAN or in the network. You configure the trust setting by using the **ip arp inspection trust** interface configuration command.

Caution Use the trust state configuration carefully. Configuring interfaces as untrusted when they should be trusted can result in a loss of connectivity.

In Figure 26-2, assume that both Switch A and Switch B are running DAI on the VLAN that includes Host 1 and Host 2. If Host 1 and Host 2 acquire their IP addresses from the DHCP server connected to Switch A, only Switch A binds the IP-to-MAC address of Host 1. Therefore, if the interface between Switch A and Switch B is untrusted, the ARP packets from Host 1 are dropped by Switch B. Connectivity between Host 1 and Host 2 is lost.

Figure 26-2 ARP Packet Validation on a VLAN Enabled for DAI

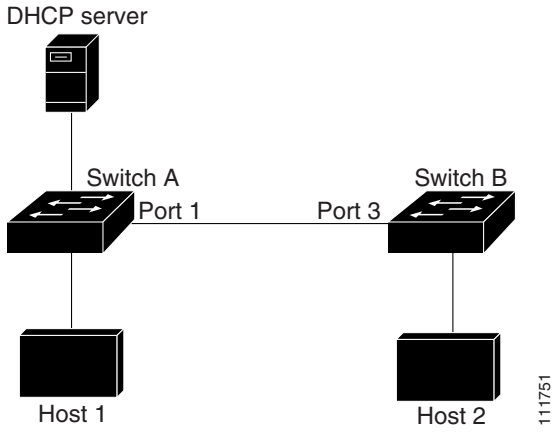

Configuring interfaces to be trusted when they are actually untrusted leaves a security hole in the network. If Switch A is not running DAI, Host 1 can easily poison the ARP cache of Switch B (and Host 2, if the link between the switches is configured as trusted). This condition can occur even though Switch B is running DAI.

DAI ensures that hosts (on untrusted interfaces) connected to a switch running DAI do not poison the ARP caches of other hosts in the network. However, DAI does not prevent hosts in other portions of the network from poisoning the caches of the hosts that are connected to a switch running DAI.

If some switches in a VLAN run DAI and other switches do not, configure the interfaces connecting these switches as untrusted. However, to validate the bindings of packets from non-DAI switches, configure the switch running DAI with ARP ACLs. When you cannot determine the bindings, at Layer 3 isolate switches running DAI from switches not running DAI switches.

Note Depending on the setup of the DHCP server and the network, it might not be possible to validate a given ARP packet on all switches in the VLAN.

Rate Limiting of ARP Packets

The switch CPU performs DAI validation checks; therefore, the number of incoming ARP packets is rate-limited to prevent a denial-of-service attack. By default, the rate for untrusted interfaces is 15 packets per second (pps). Trusted interfaces are not rate-limited. You can change this setting by using the **ip arp inspection limit** interface configuration command.

When the rate of incoming ARP packets exceeds the configured limit, the switch places the port in the error-disabled state. The port remains in that state until you intervene. You can use the **errdisable recovery** global configuration command to enable error-disable recovery so that ports automatically emerge from this state after a specified timeout period.

Note Unless you configure a rate limit on an interface, changing the trust state of the interface also changes its rate limit to the default value for that trust state. After you configure the rate limit, the interface retains the rate limit even when its trust state is changed. If you enter the **no ip arp inspection limit** interface configuration command, the interface reverts to its default rate limit.

Relative Priority of ARP ACLs and DHCP Snooping Entries

DAI uses the DHCP snooping binding database for the list of valid IP-to-MAC address bindings.

ARP ACLs take precedence over entries in the DHCP snooping binding database. The switch uses ACLs only if you configure them by using the **ip arp inspection filter vlan** global configuration command. The switch first compares ARP packets to user-configured ARP ACLs. If the ARP ACL denies the ARP packet, the switch also denies the packet even if a valid binding exists in the database populated by DHCP snooping.

Logging of Dropped Packets

When the switch drops a packet, it places an entry in the log buffer and then generates system messages on a rate-controlled basis. After the message is generated, the switch clears the entry from the log buffer. Each log entry contains flow information, such as the receiving VLAN, the port number, the source and destination IP addresses, and the source and destination MAC addresses.

You use the **ip arp inspection log-buffer** global configuration command to configure the number of entries in the buffer and the number of entries needed in the specified interval to generate system messages. You specify the type of packets that are logged by using the **ip arp inspection vlan logging** global configuration command.

A log-buffer entry can represent more than one packet. For example, if an interface receives many packets on the same VLAN with the same ARP parameters, the switch combines the packets as one entry in the log buffer and generates a single system message for the entry.

If the log buffer overflows, it means that a log event does not fit into the log buffer, and the display for the **show ip arp inspection log** privileged EXEC command is affected. Dashes in the display appears in place of all data except the packet count and the time. No other statistics are provided for the entry. If you see this entry in the display, increase the number of entries in the log buffer or increase the logging rate.

Default Dynamic ARP Inspection Settings

Table 26-1 Default Dynamic ARP Inspection Settings

Feature	Default Setting
DAI	Disabled on all VLANs.
Interface trust state	All interfaces are untrusted.
Rate limit of incoming ARP packets	The rate is 15 pps on untrusted interfaces, assuming that the network is a switched network with a host connecting to as many as 15 new hosts per second. The rate is unlimited on all trusted interfaces. The burst interval is 1 second.
ARP ACLs for non-DHCP environments	No ARP ACLs are defined.
Validation checks	No checks are performed.
Log buffer	When DAI is enabled, all denied or dropped ARP packets are logged. The number of entries in the log is 32. The number of system messages is limited to 5 per second. The logging-rate interval is 1 second.
Per-VLAN logging	All denied or dropped ARP packets are logged.

Dynamic ARP Inspection Configuration Guidelines

- DAI is an ingress security feature; it does not perform any egress checking.
- DAI is not effective for hosts connected to switches that do not support DAI or that do not have this feature enabled. Because man-in-the-middle attacks are limited to a single Layer 2 broadcast domain, separate the domain with DAI checks from the one with no checking. This action secures the ARP caches of hosts in the domain enabled for DAI.
- DAI depends on the entries in the DHCP snooping binding database to verify IP-to-MAC address bindings in incoming ARP requests and ARP responses. Make sure to enable DHCP snooping to permit ARP packets that have dynamically assigned IP addresses. For configuration information, see Chapter 25, “Configuring DHCP.”

  When DHCP snooping is disabled or in non-DHCP environments, use ARP ACLs to permit or to deny packets.
- DAI is supported on access ports, trunk ports, EtherChannel ports, and private VLAN ports.

Note Do not enable DAI on RSPAN VLANs. If DAI is enabled on RSPAN VLANs, DAI packets might not reach the RSPAN destination port.

- A physical port can join an EtherChannel port channel only when the trust state of the physical port and the channel port match. Otherwise, the physical port remains suspended in the port channel. A port channel inherits its trust state from the first physical port that joins the channel. Consequently, the trust state of the first physical port need not match the trust state of the channel.

  Conversely, when you change the trust state on the port channel, the switch configures a new trust state on all the physical ports that comprise the channel.
- The operating rate for the port channel is cumulative across all the physical ports within the channel. For example, if you configure the port channel with an ARP rate-limit of 400 pps, all the interfaces combined on the channel receive an aggregate 400 pps. The rate of incoming ARP packets on EtherChannel ports is equal to the sum of the incoming rate of packets from all the channel members. Configure the rate limit for EtherChannel ports only after examining the rate of incoming ARP packets on the channel-port members.

  The rate of incoming packets on a physical port is checked against the port-channel configuration rather than the physical-ports configuration. The rate-limit configuration on a port channel is independent of the configuration on its physical ports.

  If the EtherChannel receives more ARP packets than the configured rate, the channel (including all physical ports) is placed in the error-disabled state.
- Make sure to limit the rate of ARP packets on incoming trunk ports. Configure trunk ports with higher rates to reflect their aggregation and to handle packets across multiple DAI-enabled VLANs. You also can use the **ip arp inspection limit none** interface configuration command to make the rate unlimited. A high rate-limit on one VLAN can cause a denial-of-service attack to other VLANs when the software places the port in the error-disabled state.
- When you enable DAI on the switch, policers that were configured to police ARP traffic are no longer effective. The result is that all ARP traffic is sent to the CPU.

How to Configure Dynamic ARP Inspection

Configuring Dynamic ARP Inspection in DHCP Environments

This procedure shows how to configure DAI when two switches support this feature. Host 1 is connected to Switch A, and Host 2 is connected to Switch B as shown in Figure 26-2 on page 26-3. Both switches are running DAI on VLAN 1 where the hosts are located. A DHCP server is connected to Switch A. Both hosts acquire their IP addresses from the same DHCP server. Therefore, Switch A has the bindings for Host 1 and Host 2, and Switch B has the binding for Host 2.

Before You Begin

You must perform this procedure on both switches. This procedure is required.

	Command	Purpose
Step 1	**show cdp neighbors**	Verifies the connection between the switches.
Step 2	**configure terminal**	Enters global configuration mode.

	Command	Purpose
Step 3	**ip arp inspection vlan** *vlan-range*	Enables DAI on a per-VLAN basis. By default, DAI is disabled on all VLANs. *vlan-range*—Specifies a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. The range is 1 to 4096. Specifies the same VLAN ID for both switches.
Step 4	**interface** *interface-id*	Specifies the interface connected to the other switch, and enters interface configuration mode.
Step 5	**ip arp inspection trust**	Configures the connection between the switches as trusted. By default, all interfaces are untrusted. The switch does not check ARP packets that it receives from the other switch on the trusted interface; it only forwards the packets. For untrusted interfaces, the switch intercepts all ARP requests and responses. It verifies that the intercepted packets have valid IP-to-MAC address bindings before updating the local cache and before forwarding the packet to the appropriate destination. The switch drops invalid packets and logs them in the log buffer according to the logging configuration specified with the **ip arp inspection vlan logging** global configuration command.
Step 6	**end**	Returns to privileged EXEC mode.

Configuring ARP ACLs for Non-DHCP Environments

This procedure shows how to configure DAI when Switch B shown in Figure 26-2 on page 26-3 does not support DAI or DHCP snooping.

If you configure port 1 on Switch A as trusted, a security hole is created because both Switch A and Host 1 could be attacked by either Switch B or Host 2. To prevent this possibility, you must configure port 1 on Switch A as untrusted. To permit ARP packets from Host 2, you must set up an ARP ACL and apply it to VLAN 1. If the IP address of Host 2 is not static (it is impossible to apply the ACL configuration on Switch A) you must separate Switch A from Switch B at Layer 3 and use a router to route packets between them.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**arp access-list** *acl-name*	Defines an ARP ACL, and enters ARP access-list configuration mode. By default, no ARP access lists are defined. **Note** At the end of the ARP access list, there is an implicit **deny ip any mac any** command.

	Command	Purpose
Step 3	**permit ip host** *sender-ip* **mac host** *sender-mac* [**log**]	Permits ARP packets from the specified host (Host 2). • *sender-ip*—Enters the IP address of Host 2. • *sender-mac*—Enters the MAC address of Host 2. • (Optional) **log**—Logs a packet in the log buffer when it matches the access control entry (ACE). Matches are logged if you also configure the **matchlog** keyword in the **ip arp inspection vlan logging** global configuration command. For more information, see the "Configuring the Log Buffer" section on page 26-11.
Step 4	**exit**	Returns to global configuration mode.
Step 5	**ip arp inspection filter** *arp-acl-name* **vlan** *vlan-range* [**static**]	Applies the ARP ACL to the VLAN. By default, no defined ARP ACLs are applied to any VLAN. • *arp-acl-name*—Specifies the name of the ACL created in Step 2. • *vlan-range*—Specifies the VLAN that the switches and hosts are in. You can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. The range is 1 to 4096. • (Optional) **static**—Specifies to treat implicit denies in the ARP ACL as explicit denies and to drop packets that do not match any previous clauses in the ACL. DHCP bindings are not used. If you do not specify this keyword, it means that there is no explicit deny in the ACL that denies the packet, and DHCP bindings determine whether a packet is permitted or denied if the packet does not match any clauses in the ACL. ARP packets containing only IP-to-MAC address bindings are compared against the ACL. Packets are permitted only if the access list permits them.
Step 6	**interface** *interface-id*	Specifies the Switch A interface that is connected to Switch B, and enters interface configuration mode.
Step 7	**no ip arp inspection trust**	Configures the Switch A interface that is connected to Switch B as untrusted. By default, all interfaces are untrusted. For untrusted interfaces, the switch intercepts all ARP requests and responses. It verifies that the intercepted packets have valid IP-to-MAC address bindings before updating the local cache and before forwarding the packet to the appropriate destination. The switch drops invalid packets and logs them in the log buffer according to the logging configuration specified with the **ip arp inspection vlan logging** global configuration command.
Step 8	**end**	Returns to privileged EXEC mode.

Limiting the Rate of Incoming ARP Packets

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface to be rate-limited, and enters interface configuration mode.
Step 3	**ip arp inspection limit** {**rate** *pps* [**burst interval** *seconds*] | **none**}	Limits the rate of incoming ARP requests and responses on the interface. The default rate is 15 pps on untrusted interfaces and unlimited on trusted interfaces. The burst interval is 1 second. • **rate** *pps*—Specifies an upper limit for the number of incoming packets processed per second. The range is 0 to 2048 pps. • (Optional) **burst interval** *seconds*—Specifies the consecutive interval in seconds, over which the interface is monitored for a high rate of ARP packets. The range is 1 to 15. • **rate none**—Specifies no upper limit for the rate of incoming ARP packets that can be processed.
Step 4	**exit**	Returns to global configuration mode.
Step 5	**errdisable recovery cause arp-inspection interval** *interval*	(Optional) Enables error recovery from the DAI error-disabled state. By default, recovery is disabled, and the recovery interval is 300 seconds. **interval** *interval*—Specifies the time in seconds to recover from the error-disabled state. The range is 30 to 86400.
Step 6	**exit**	Returns to privileged EXEC mode.

Performing Validation Checks

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip arp inspection validate** {[**src-mac**] [**dst-mac**] [**ip**]}	Performs a specific check on incoming ARP packets. By default, no checks are performed. • **src-mac**—Checks the source MAC address in the Ethernet header against the sender MAC address in the ARP body. This check is performed on both ARP requests and responses. When enabled, packets with different MAC addresses are classified as invalid and are dropped. • **dst-mac**—Checks the destination MAC address in the Ethernet header against the target MAC address in ARP body. This check is performed for ARP responses. When enabled, packets with different MAC addresses are classified as invalid and are dropped. • **ip**—Checks the ARP body for invalid and unexpected IP addresses. Addresses include 0.0.0.0, 255.255.255.255, and all IP multicast addresses. Sender IP addresses are checked in all ARP requests and responses, and target IP addresses are checked only in ARP responses. You must specify at least one of the keywords. Each command overrides the configuration of the previous command; that is, if a command enables **src** and **dst mac** validations, and a second command enables IP validation only, the **src** and **dst mac** validations are disabled as a result of the second command.
Step 3	**exit**	Returns to privileged EXEC mode.

Configuring the Log Buffer

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip arp inspection log-buffer** {**entries** *number* | **logs** *number* **interval** *seconds*}	Configures the DAI logging buffer. By default, when DAI is enabled, denied, or dropped, ARP packets are logged. The number of log entries is 32. The number of system messages is limited to 5 per second. The logging-rate interval is 1 second. • **entries** *number*—Specifies the number of entries to be logged in the buffer. The range is 0 to 1024. • **logs** *number* **interval** *seconds*—Specifies the number of entries to generate system messages in the specified interval. **logs** *number*—Specifies the range 0 to 1024. A 0 value means that the entry is placed in the log buffer, but a system message is not generated. **interval** *seconds*—Specifies the range 0 to 86400 seconds (1 day). A 0 value means that a system message is immediately generated (and the log buffer is always empty). An interval setting of 0 overrides a log setting of 0. The **logs** and **interval** settings interact. If the **logs** *number* X is greater than **interval** *seconds* Y, X divided by Y (X/Y) system messages are sent every second. Otherwise, one system message is sent every Y divided by X (Y/X) seconds.
Step 3	**ip arp inspection vlan** *vlan-range* **logging** {**acl-match** {**matchlog** | **none**} | **dhcp-bindings** {**all** | **none** | **permit**}}	Controls the type of packets that are logged per VLAN. By default, all denied or all dropped packets are logged. The term *logged* means the entry is placed in the log buffer and a system message is generated. • *vlan-range*—Specifies a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. The range is 1 to 4096. • **acl-match matchlog**—Specifies log packets based on the ACE logging configuration. If you specify the **matchlog** keyword in this command and the **log** keyword in the **permit** or **deny** ARP access-list configuration command, ARP packets permitted or denied by the ACL are logged. • **acl-match none**—Does not log packets that match ACLs. • **dhcp-bindings all**—Logs all packets that match DHCP bindings. • **dhcp-bindings none**—Does not log packets that match DHCP bindings. • **dhcp-bindings permit**—Logs DHCP-binding permitted packets.
Step 4	**exit**	Returns to privileged EXEC mode.

Monitoring and Maintaining Dynamic ARP Inspection

Command	Description
clear ip arp inspection log	Clears the DAI log buffer.
clear ip arp inspection statistics	Clears the DAI statistics.
show arp access-list [*acl-name*]	Displays detailed information about ARP ACLs.
show errdisable recovery	Displays the error-disabled recovery timer information.
show ip arp inspection interfaces [*interface-id*]	Displays the trust state and the rate limit of ARP packets for the specified interface or all interfaces.
show ip arp inspection log	Displays the configuration and contents of the DAI log buffer.
show ip arp inspection vlan *vlan-range*	Displays the configuration and the operating state of DAI for the specified VLAN. If no VLANs are specified or if a range is specified, displays information only for VLANs with DAI enabled (active).
show ip arp inspection statistics [**vlan** *vlan-range*]	Displays statistics for forwarded, dropped, MAC validation failure, IP validation failure, ACL permitted and denied, and DHCP permitted and denied packets for the specified VLAN. If no VLANs are specified or if a range is specified, displays information only for VLANs with DAI enabled (active).
show ip dhcp snooping binding	Verifies the DHCP bindings.

Configuration Examples for Dynamic ARP Inspection

Configuring Dynamic ARP Inspection in DHCP Environments: Example

This example shows how to configure DAI on Switch A in VLAN 1. You would perform a similar procedure on Switch B:

```
Switch(config)# ip arp inspection vlan 1
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# ip arp inspection trust
```

Configuring ARP ACLs for Non-DHCP Environments: Example

This example shows how to configure an ARP ACL called host2 on Switch A, to permit ARP packets from Host 2 (IP address 1.1.1.1 and MAC address 0001.0001.0001), to apply the ACL to VLAN 1, and to configure port 1 on Switch A as untrusted:

```
Switch(config)# arp access-list host2
Switch(config-arp-acl)# permit ip host 1.1.1.1 mac host 1.1.1
Switch(config-arp-acl)# exit
Switch(config)# ip arp inspection filter host2 vlan 1
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# no ip arp inspection trust
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
DHCP configuration	"Configuring DHCP on the IE 2000 Switch"

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 27

Configuring IP Source Guard

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for IP Source Guard

- You must globally configure the **ip device tracking maximum** *limit-number* interface configuration command globally for IPSG for static hosts to work. If you only configure this command on a port without enabling IP device tracking globally or setting an IP device tracking maximum on that interface, IPSG with static hosts will reject all the IP traffic from that interface. This requirement also applies to IPSG with static hosts on a Layer 2 access port.

Restrictions for IP Source Guard

- To use this feature, the switch must be running the LAN Base image.
- IP source guard (IPSG) is supported only on Layer 2 ports, including access and trunk ports.
- Do not use IPSG for static hosts on uplink ports or trunk ports.

Information About IP Source Guard

IP Source Guard

IPSG is a security feature that restricts IP traffic on nonrouted, Layer 2 interfaces by filtering traffic based on the DHCP snooping binding database and on manually configured IP source bindings. You can use IPSG to prevent traffic attacks if a host tries to use the IP address of its neighbor.

You can enable IPSG when DHCP snooping is enabled on an untrusted interface. After IPSG is enabled on an interface, the switch blocks all IP traffic received on the interface except for DHCP packets allowed by DHCP snooping. A port access control list (ACL) is applied to the interface. The port ACL allows only IP traffic with a source IP address in the IP source binding table and denies all other traffic.

Note The port ACL takes precedence over any router ACLs or VLAN maps that affect the same interface.

The IP source binding table bindings are learned by DHCP snooping or are manually configured (static IP source bindings). An entry in this table has an IP address with its associated MAC address and VLAN number. The switch uses the IP source binding table only when IPSG is enabled.

You can configure IPSG with source IP address filtering or with source IP and MAC address filtering.

Source IP Address Filtering

When IPSG is enabled with this option, IP traffic is filtered based on the source IP address. The switch forwards IP traffic when the source IP address matches an entry in the DHCP snooping binding database or a binding in the IP source binding table.

When a DHCP snooping binding or static IP source binding is added, changed, or deleted on an interface, the switch modifies the port ACL by using the IP source binding changes and re-applies the port ACL to the interface.

If you enable IPSG on an interface on which IP source bindings (dynamically learned by DHCP snooping or manually configured) are not configured, the switch creates and applies a port ACL that denies all IP traffic on the interface. If you disable IPSG, the switch removes the port ACL from the interface.

Source IP and MAC Address Filtering

IP traffic is filtered based on the source IP and MAC addresses. The switch forwards traffic only when the source IP and MAC addresses match an entry in the IP source binding table.

When address filtering is enabled, the switch filters IP and non-IP traffic. If the source MAC address of an IP or non-IP packet matches a valid IP source binding, the switch forwards the packet. The switch drops all other types of packets except DHCP packets.

The switch uses port security to filter source MAC addresses. The interface can shut down when a port-security violation occurs.

IP Source Guard for Static Hosts

IPSG for static hosts extends the IPSG capability to non-DHCP and static environments. The previous IPSG used the entries created by DHCP snooping to validate the hosts connected to a switch. Any traffic received from a host without a valid DHCP binding entry is dropped. This security feature restricts IP traffic on nonrouted Layer 2 interfaces. It filters traffic based on the DHCP snooping binding database and on manually configured IP source bindings. The previous version of IPSG required a DHCP environment for IPSG to work.

IPSG for static hosts allows IPSG to work without DHCP. IPSG for static hosts relies on IP device tracking-table entries to install port ACLs. The switch creates static entries based on ARP requests or other IP packets to maintain the list of valid hosts for a given port. You can also specify the number of hosts allowed to send traffic to a given port. This is equivalent to port security at Layer 3.

IPSG for static hosts also supports dynamic hosts. If a dynamic host receives a DHCP-assigned IP address that is available in the IP DHCP snooping table, the same entry is learned by the IP device tracking table. When you enter the **show ip device tracking all** EXEC command, the IP device tracking table displays the entries as ACTIVE.

Note Some IP hosts with multiple network interfaces can inject some invalid packets into a network interface. The invalid packets contain the IP or MAC address for another network interface of the host as the source address. The invalid packets can cause IPSG for static hosts to connect to the host, to learn the invalid IP or MAC address bindings, and to reject the valid bindings. Consult the vendor of the corresponding operating system and the network interface to prevent the host from injecting invalid packets.

IPSG for static hosts initially learns IP or MAC bindings dynamically through an ACL-based snooping mechanism. IP or MAC bindings are learned from static hosts by ARP and IP packets. They are stored in the device tracking database. When the number of IP addresses that have been dynamically learned or statically configured on a given port reaches a maximum, the hardware drops any packet with a new IP address. To resolve hosts that have moved or gone away for any reason, IPSG for static hosts leverages IP device tracking to age out dynamically learned IP address bindings. This feature can be used with DHCP snooping. Multiple bindings are established on a port that is connected to both DHCP and static hosts. For example, bindings are stored in both the device tracking database as well as in the DHCP snooping binding database.

IP Source Guard Configuration Guidelines

- By default, IP source guard is disabled.
- You can configure static IP bindings only on nonrouted ports. If you enter the **ip source binding** *mac-address* **vlan** *vlan-id ip-address* **interface** *interface-id* global configuration command on a routed interface, this error message appears:

  ```
  Static IP source binding can only be configured on switch port.
  ```

- When IP source guard with source IP filtering is enabled on an interface, DHCP snooping must be enabled on the access VLAN for that interface.
- If you are enabling IP source guard on a trunk interface with multiple VLANs and DHCP snooping is enabled on all the VLANs, the source IP address filter is applied on all the VLANs.

  Note If IP source guard is enabled and you enable or disable DHCP snooping on a VLAN on the trunk interface, the switch might not properly filter traffic.

- If you enable IP source guard with source IP and MAC address filtering, DHCP snooping and port security must be enabled on the interface. You must also enter the **ip dhcp snooping information option** global configuration command and ensure that the DHCP server supports option 82. When IP source guard is enabled with MAC address filtering, the DHCP host MAC address is not learned until the host is granted a lease. When forwarding packets from the server to the host, DHCP snooping uses option-82 data to identify the host port.
- When configuring IP source guard on interfaces on which a private VLAN is configured, port security is not supported.
- IP source guard is not supported on EtherChannels.

- You can enable this feature when 802.1x port-based authentication is enabled.
- If the number of ternary content addressable memory (TCAM) entries exceeds the maximum, the CPU usage increases.

How to Configure IP Source Guard

Enabling IP Source Guard

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface to be configured, and enters interface configuration mode.
Step 3	**ip verify source** or **ip verify source port-security**	Enables IPSG with source IP address filtering. Enables IPSG with source IP and MAC address filtering. **Note** When you enable both IPSG and port security by using the **ip verify source port-security** interface configuration command, there are two caveats: • The DHCP server must support option-82, or the client is not assigned an IP address. • The MAC address in the DHCP packet is not learned as a secure address. The MAC address of the DHCP client is learned as a secure address only when the switch receives non-DHCP data traffic.
Step 4	**exit**	Returns to global configuration mode.
Step 5	**ip source binding** *mac-address* **vlan** *vlan-id ip-address* **inteface** *interface-id*	Adds a static IP source binding. Enter this command for each static binding.
Step 6	**end**	Returns to privileged EXEC mode.

Configuring IP Source Guard for Static Hosts on a Layer 2 Access Port

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip device tracking**	Opens the IP host table, and globally enables IP device tracking.
Step 3	**interface** *interface-id*	Enters interface configuration mode.
Step 4	**switchport mode access**	Configures a port as access.
Step 5	**switchport access vlan** *vlan-id*	Configures the VLAN for this port.

	Command	Purpose
Step 6	**ip verify source tracking port-security**	Enables IPSG for static hosts with MAC address filtering. **Note** When you enable both IPSG and port security by using the **ip verify source port-security** interface configuration command: • The DHCP server must support option-82, or the client is not assigned an IP address. • The MAC address in the DHCP packet is not learned as a secure address. The MAC address of the DHCP client is learned as a secure address only when the switch receives non-DHCP data traffic.
Step 7	**ip device tracking maximum** *number*	Specifies a maximum limit for the number of static IPs that the IP device tracking table allows on the port. The range is 1to 10. The maximum number is 10. **Note** You must configure the **ip device tracking maximum** *limit-number* interface configuration command.
Step 8	**switchport port-security**	(Optional) Activates port security for this port.
Step 9	**switchport port-security maximum** *value*	(Optional) Specifies a maximum of MAC addresses for this port.
Step 10	**end**	Returns to privileged EXEC mode.
Step 11	**show ip verify source interface** *interface-id*	Verifies the configuration and displays IPSG permit ACLs for static hosts.
Step 12	**show ip device track all [active | inactive] count**	Verifies the configuration by displaying the IP-to-MAC binding for a given host on the switch interface. • **all active**—Displays only the active IP or MAC binding entries • **all inactive**—Displays only the inactive IP or MAC binding entries • **all**—Displays the active and inactive IP or MAC binding entries

Configuring IP Source Guard for Static Hosts on a Private VLAN Host Port

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**vlan** *vlan-id1*	Enters VLAN configuration mode.
Step 3	**private-vlan primary**	Specifies a primary VLAN on a private VLAN port.
Step 4	**exit**	Exits VLAN configuration mode.
Step 5	**vlan** *vlan-id2*	Enters configuration VLAN mode for another VLAN.

	Command	Purpose
Step 6	**private-vlan isolated**	Specifies an isolated VLAN on a private VLAN port.
Step 7	**exit**	Exits VLAN configuration mode.
Step 8	**vlan** *vlan-id1*	Enters configuration VLAN mode.
Step 9	**private-vlan association 201**	Associates the VLAN on an isolated private VLAN port.
Step 10	**exit**	Exits VLAN configuration mode.
Step 11	**interface fastEthernet** *interface-id*	Enters interface configuration mode.
Step 12	**switchport mode private-vlan host**	(Optional) Specifies a port as a private VLAN host.
Step 13	**switchport private-vlan host-association** *vlan-id1 vlan-id2*	(Optional) Associates this port with the corresponding private VLAN.
Step 14	**ip device tracking maximum** *number*	Specifies a maximum for the number of static IPs that the IP device tracking table allows on the port. The maximum is 10. **Note** You must globally configure the **ip device tracking maximum** *number* interface command for IPSG for static hosts to work.
Step 15	**ip verify source tracking [port-security]**	Activates IPSG for static hosts with MAC address filtering on this port.
Step 16	**end**	Exits configuration interface mode.
Step 17	**show ip device tracking all**	Verifies the configuration.
Step 18	**show ip verify source interface** *interface-id*	Verifies the IPSG configuration and displays IPSG permit ACLs for static hosts.

Monitoring and Maintaining IP Source Guard

Command	Purpose
show ip device tracking	Displays the active IP or MAC binding entries for all interfaces.
show ip source binding	Displays the IP source bindings on a switch.
show ip verify source	Displays the IP source guard configuration on the switch.
copy running-config startup-config	Saves your entries in the configuration file.

Configuration Examples for IP Source Guard

Enabling IPSG with Source IP and MAC Filtering: Example

This example shows how to enable IPSG with source IP and MAC filtering on VLANs 10 and 11:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# ip verify source port-security
Switch(config-if)# exit
Switch(config)# ip source binding 0100.0022.0010 vlan 10 10.0.0.2 interface
gigabitethernet1/1
Switch(config)# ip source binding 0100.0230.0002 vlan 11 10.0.0.4 interface
gigabitethernet1/1
Switch(config)# end
```

Disabling IPSG with Static Hosts: Example

This example shows how to stop IPSG with static hosts on an interface:

```
Switch(config-if)# no ip verify source
Switch(config-if)# no ip device tracking max
```

Enabling IPSG for Static Hosts: Examples

This example shows how to enable IPSG with static hosts on a port:

```
Switch(config)# ip device tracking
Switch(config)# ip device tracking max 10
Switch(config-if)# ip verify source tracking port-security
```

This example shows how to enable IPSG for static hosts with IP filters on a Layer 2 access port and to verify the valid IP bindings on the interface Gi0/3:

```
Switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# ip device tracking
Switch(config)# interface gigabitethernet 0/3
Switch(config-if)# switchport mode access
```

```
Switch(config-if)# switchport access vlan 10
Switch(config-if)# ip device tracking maximum 5
Switch(config-if)# ip verify source tracking
Switch(config-if)# end

Switch# show ip verify source
Interface  Filter-type  Filter-mode  IP-address       Mac-address        Vlan
---------  -----------  -----------  ---------------  -----------------  ----
Gi0/3      ip trk       active       40.1.1.24                           10
Gi0/3      ip trk       active       40.1.1.20                           10
Gi0/3      ip trk       active       40.1.1.21                           10
```

This example shows how to enable IPSG for static hosts with IP-MAC filters on a Layer 2 access port, to verify the valid IP-MAC bindings on the interface Gi0/3, and to verify that the number of bindings on this interface has reached the maximum:

```
Switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# ip device tracking
Switch(config)# interface gigabitethernet 0/3
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan 1
Switch(config-if)# ip device tracking maximum 5
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security maximum 5
Switch(config-if)# ip verify source tracking port-security
Switch(config-if)# end

Switch# show ip verify source
Interface  Filter-type  Filter-mode  IP-address       Mac-address        Vlan
---------  -----------  -----------  ---------------  -----------------  ----
Gi0/3      ip-mac trk   active       40.1.1.24        00:00:00:00:03:04  1
Gi0/3      ip-mac trk   active       40.1.1.20        00:00:00:00:03:05  1
Gi0/3      ip-mac trk   active       40.1.1.21        00:00:00:00:03:06  1
Gi0/3      ip-mac trk   active       40.1.1.22        00:00:00:00:03:07  1
Gi0/3      ip-mac trk   active       40.1.1.23        00:00:00:00:03:08  1
```

Displaying IP or MAC Binding Entries: Examples

This example displays all IP or MAC binding entries for all interfaces. The CLI displays all active as well as inactive entries. When a host is learned on a interface, the new entry is marked as active. When the same host is disconnected from that interface and connected to a different interface, a new IP or MAC binding entry displays as active as soon as the host is detected. The old entry for this host on the previous interface is marked as INACTIVE.

```
Switch# show ip device tracking all
IP Device Tracking = Enabled
IP Device Tracking Probe Count = 3
IP Device Tracking Probe Interval = 30
-----------------------------------------------------------------------
  IP Address     MAC Address   Vlan  Interface              STATE
-----------------------------------------------------------------------
200.1.1.8       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
200.1.1.9       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
200.1.1.10      0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
200.1.1.1       0001.0600.0000  9    GigabitEthernet0/2     ACTIVE
200.1.1.1       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
200.1.1.2       0001.0600.0000  9    GigabitEthernet0/2     ACTIVE
200.1.1.2       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
200.1.1.3       0001.0600.0000  9    GigabitEthernet0/2     ACTIVE
200.1.1.3       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
```

```
200.1.1.4       0001.0600.0000  9    GigabitEthernet0/2     ACTIVE
200.1.1.4       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
200.1.1.5       0001.0600.0000  9    GigabitEthernet0/2     ACTIVE
200.1.1.5       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
200.1.1.6       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
200.1.1.7       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
```

This example displays all active IP or MAC binding entries for all interfaces:

```
Switch# show ip device tracking all active
IP Device Tracking = Enabled
IP Device Tracking Probe Count = 3
IP Device Tracking Probe Interval = 30
-----------------------------------------------------------------------
  IP Address     MAC Address   Vlan  Interface              STATE
-----------------------------------------------------------------------
200.1.1.1       0001.0600.0000  9    GigabitEthernet0/1     ACTIVE
200.1.1.2       0001.0600.0000  9    GigabitEthernet0/1     ACTIVE
200.1.1.3       0001.0600.0000  9    GigabitEthernet0/1     ACTIVE
200.1.1.4       0001.0600.0000  9    GigabitEthernet0/1     ACTIVE
200.1.1.5       0001.0600.0000  9    GigabitEthernet0/1     ACTIVE
```

This example displays all inactive IP or MAC binding entries for all interfaces. The host was first learned on GigabitEthernet 0/1 and then moved to GigabitEthernet 0/2. The IP or MAC binding entries learned on GigabitEthernet 0/1 are marked as inactive.

```
Switch# show ip device tracking all inactive
IP Device Tracking = Enabled
IP Device Tracking Probe Count = 3
IP Device Tracking Probe Interval = 30
-----------------------------------------------------------------------
  IP Address     MAC Address   Vlan  Interface              STATE
-----------------------------------------------------------------------
200.1.1.8       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
200.1.1.9       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
200.1.1.10      0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
200.1.1.1       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
200.1.1.2       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
200.1.1.3       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
200.1.1.4       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
200.1.1.5       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
200.1.1.6       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
200.1.1.7       0001.0600.0000  8    GigabitEthernet0/1     INACTIVE
```

This example displays the count of all IP device tracking host entries for all interfaces:

```
Switch# show ip device tracking all count
Total IP Device Tracking Host entries: 5
-----------------------------------------------------------------------
  Interface           Maximum Limit           Number of Entries
-----------------------------------------------------------------------
Gi0/3                  5
```

Enabling IPSG for Static Hosts: Examples

This example shows how to enable IPSG for static hosts with IP filters on a private VLAN host port:

```
Switch(config)# vlan 200
Switch(config-vlan)# private-vlan primary
Switch(config-vlan)# exit
Switch(config)# vlan 201
Switch(config-vlan)# private-vlan isolated
```

```
Switch(config-vlan)# exit
Switch(config)# vlan 200
Switch(config-vlan)# private-vlan association 201
Switch(config-vlan)# exit
Switch(config)# int fastEthernet 4/3
Switch(config-if)# switchport mode private-vlan host
Switch(config-if)# switchport private-vlan host-association 200 201
Switch(config-if)# ip device tracking maximum 8
Switch(config-if)# ip verify source tracking

Switch# show ip device tracking all
IP Device Tracking = Enabled
IP Device Tracking Probe Count = 3
IP Device Tracking Probe Interval = 30
-----------------------------------------------------------------
  IP Address     MAC Address   Vlan  Interface              STATE
-----------------------------------------------------------------
40.1.1.24       0000.0000.0304  200  FastEthernet0/3        ACTIVE
40.1.1.20       0000.0000.0305  200  FastEthernet0/3        ACTIVE
40.1.1.21       0000.0000.0306  200  FastEthernet0/3        ACTIVE
40.1.1.22       0000.0000.0307  200  FastEthernet0/3        ACTIVE
40.1.1.23       0000.0000.0308  200  FastEthernet0/3        ACTIVE
```

The output shows the five valid IP-MAC bindings that have been learned on the interface Fa0/3. For the private VLAN cases, the bindings are associated with primary VLAN ID. In this example, the primary VLAN ID, 200, is shown in the table.

```
Switch# show ip verify source
Interface  Filter-type  Filter-mode  IP-address       Mac-address        Vlan
---------  -----------  -----------  ---------------  -----------------  ----
Fa0/3      ip trk       active       40.1.1.23                           200
Fa0/3      ip trk       active       40.1.1.24                           200
Fa0/3      ip trk       active       40.1.1.20                           200
Fa0/3      ip trk       active       40.1.1.21                           200
Fa0/3      ip trk       active       40.1.1.22                           200
Fa0/3      ip trk       active       40.1.1.23                           201
Fa0/3      ip trk       active       40.1.1.24                           201
Fa0/3      ip trk       active       40.1.1.20                           201
Fa0/3      ip trk       active       40.1.1.21                           201
Fa0/30/3   ip trk       active       40.1.1.22                           201
```

The output shows that the five valid IP-MAC bindings are on both the primary and secondary VLAN.

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

CHAPTER 28

Configuring IGMP Snooping and MVR

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Restrictions for IGMP Snooping and MVR

- To use the Multicast VLAN Registration (MVR) feature, the switch must be running the LAN Base image.
- You can set the maximum number of IGMP groups that a Layer 2 interface can join by using the **ip igmp max-groups** interface configuration command. Use the **no** form of this command to set the maximum back to the default, which is no limit. This restriction can be applied to Layer 2 ports only—you cannot set a maximum number of IGMP groups on routed ports or SVIs. You also can use this command on a logical EtherChannel interface but cannot use it on ports that belong to an EtherChannel port group.

Information About IGMP Snooping and MVR

This chapter describes how to configure Internet Group Management Protocol (IGMP) snooping on the switch, including an application of local IGMP snooping, Multicast VLAN Registration (MVR). It also includes procedures for controlling multicast group membership by using IGMP filtering and procedures for configuring the IGMP throttling action.

Note For IP Version 6 (IPv6) traffic, Multicast Listener Discovery (MLD) snooping performs the same function as IGMP snooping for IPv4 traffic. For information about MLD snooping, see Chapter 44, "Configuring IPv6 MLD Snooping."

Note You can either manage IP multicast group addresses through features such as IGMP snooping and MVR, or you can use static IP addresses.

IGMP Snooping

Layer 2 switches can use IGMP snooping to constrain the flooding of multicast traffic by dynamically configuring Layer 2 interfaces so that multicast traffic is forwarded to only those interfaces associated with IP multicast devices. As the name implies, IGMP snooping requires the LAN switch to snoop on the IGMP transmissions between the host and the router and to keep track of multicast groups and member ports. When the switch receives an IGMP report from a host for a particular multicast group, the switch adds the host port number to the forwarding table entry; when it receives an IGMP Leave Group message from a host, it removes the host port from the table entry. It also periodically deletes entries if it does not receive IGMP membership reports from the multicast clients.

Note For more information on IP multicast and IGMP, see RFC 1112 and RFC 2236.

The multicast router sends out periodic general queries to all VLANs. All hosts interested in this multicast traffic send join requests and are added to the forwarding table entry. The switch creates one entry per VLAN in the IGMP snooping IP multicast forwarding table for each group from which it receives an IGMP join request.

The switch supports IP multicast group-based bridging, rather than MAC-addressed based groups. With multicast MAC address-based groups, if an IP address being configured translates (aliases) to a previously configured MAC address or to any reserved multicast MAC addresses (in the range 224.0.0.xxx), the command fails. Because the switch uses IP multicast groups, there are no address aliasing issues.

The IP multicast groups learned through IGMP snooping are dynamic. However, you can statically configure multicast groups by using the **ip igmp snooping vlan** *vlan-id* **static** *ip_address* **interface** *interface-id* global configuration command. If you specify group membership for a multicast group address statically, your setting supersedes any automatic manipulation by IGMP snooping. Multicast group membership lists can consist of both user-defined and IGMP snooping-learned settings.

You can configure an IGMP snooping querier to support IGMP snooping in subnets without multicast interfaces because the multicast traffic does not need to be routed. For more information about the IGMP snooping querier, see the "Configuring the IGMP Snooping Querier" section on page 28-16.

If a port spanning-tree, a port group, or a VLAN ID change occurs, the IGMP snooping-learned multicast groups from this port on the VLAN are deleted.

When you enable IGMP Immediate Leave, the switch immediately removes a port when it detects an IGMP Version 2 leave message on that port. You should only use the Immediate-Leave feature when there is a single receiver present on every port in the VLAN.

IGMP Versions

The switch supports IGMP Version 1, IGMP Version 2, and IGMP Version 3. These versions are interoperable on the switch. For example, if IGMP snooping is enabled on an IGMPv2 switch and the switch receives an IGMPv3 report from a host, the switch can forward the IGMPv3 report to the multicast router.

Note The switch supports IGMPv3 snooping based only on the destination multicast MAC address. It does not support snooping based on the source MAC address or on proxy reports.

An IGMPv3 switch supports Basic IGMPv3 Snooping Support (BISS), which includes support for the snooping features on IGMPv1 and IGMPv2 switches and for IGMPv3 membership report messages. BISS constrains the flooding of multicast traffic when your network includes IGMPv3 hosts. It constrains traffic to approximately the same set of ports as the IGMP snooping feature on IGMPv2 or IGMPv1 hosts.

Note IGMPv3 join and leave messages are not supported on switches running IGMP filtering or MVR.

An IGMPv3 switch can receive messages from and forward messages to a device running the Source Specific Multicast (SSM) feature.

Joining a Multicast Group

When a host connected to the switch wants to join an IP multicast group and it is an IGMP Version 2 client, it sends an unsolicited IGMP join message, specifying the IP multicast group to join. Alternatively, when the switch receives a general query from the router, it forwards the query to all ports in the VLAN. IGMP Version 1 or Version 2 hosts wanting to join the multicast group respond by sending a join message to the switch. The switch CPU creates a multicast forwarding-table entry for the group if it is not already present. The CPU also adds the interface where the join message was received to the forwarding-table entry. The host associated with that interface receives multicast traffic for that multicast group. See Figure 28-1.

Figure 28-1 Initial IGMP Join Message

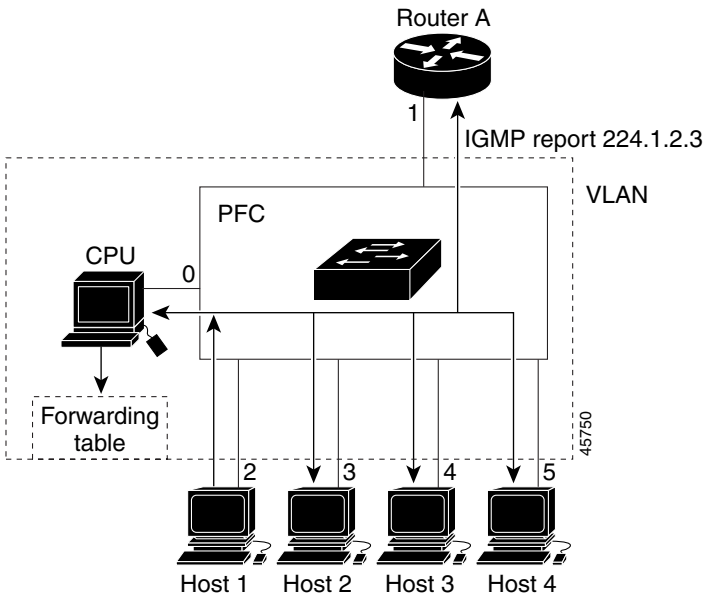

Router A sends a general query to the switch, which forwards the query to ports 2 through 5, which are all members of the same VLAN. Host 1 wants to join multicast group 224.1.2.3 and multicasts an IGMP membership report (IGMP join message) to the group. The switch CPU uses the information in the IGMP report to set up a forwarding-table entry, as shown in Table 28-1, that includes the port numbers connected to Host 1 and the router.

Table 28-1 IGMP Snooping Forwarding Table

Destination Address	Type of Packet	Ports
224.1.2.3	IGMP	1, 2

The switch hardware can distinguish IGMP information packets from other packets for the multicast group. The information in the table tells the switching engine to send frames addressed to the 224.1.2.3 multicast IP address that are not IGMP packets to the router and to the host that has joined the group.

If another host (for example, Host 4) sends an unsolicited IGMP join message for the same group (Figure 28-2), the CPU receives that message and adds the port number of Host 4 to the forwarding table as shown in Table 28-2. Note that because the forwarding table directs IGMP messages only to the CPU, the message is not flooded to other ports on the switch. Any known multicast traffic is forwarded to the group and not to the CPU.

Figure 28-2 Second Host Joining a Multicast Group

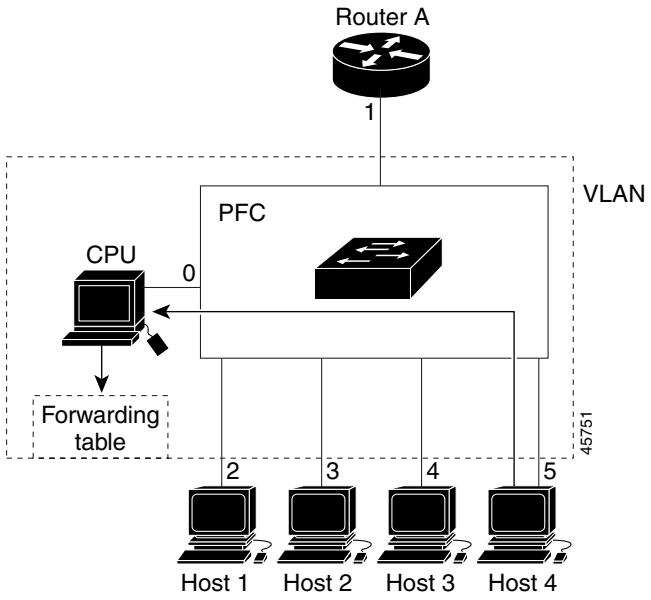

Table 28-2 Updated IGMP Snooping Forwarding Table

Destination Address	Type of Packet	Ports
224.1.2.3	IGMP	1, 2, 5

Leaving a Multicast Group

The router sends periodic multicast general queries, and the switch forwards these queries through all ports in the VLAN. Interested hosts respond to the queries. If at least one host in the VLAN wishes to receive multicast traffic, the router continues forwarding the multicast traffic to the VLAN. The switch forwards multicast group traffic only to those hosts listed in the forwarding table for that IP multicast group maintained by IGMP snooping.

When hosts want to leave a multicast group, they can silently leave, or they can send a leave message. When the switch receives a leave message from a host, it sends a group-specific query to learn if any other devices connected to that interface are interested in traffic for the specific multicast group. The switch then updates the forwarding table for that MAC group so that only those hosts interested in receiving multicast traffic for the group are listed in the forwarding table. If the router receives no reports from a VLAN, it removes the group for the VLAN from its IGMP cache.

Immediate Leave

Immediate Leave is only supported on IGMP Version 2 hosts.

The switch uses IGMP snooping Immediate Leave to remove from the forwarding table an interface that sends a leave message without the switch sending group-specific queries to the interface. The VLAN interface is pruned from the multicast tree for the multicast group specified in the original leave message. Immediate Leave ensures optimal bandwidth management for all hosts on a switched network, even when multiple multicast groups are simultaneously in use.

Note You should only use the Immediate Leave feature on VLANs where a single host is connected to each port. If Immediate Leave is enabled in VLANs where more than one host is connected to a port, some hosts might inadvertently be dropped.

When you enable IGMP Immediate Leave, the switch immediately removes a port when it detects an IGMP Version 2 leave message on that port. You should only use the Immediate-Leave feature when there is a single receiver present on every port in the VLAN.

IGMP Configurable-Leave Timer

You can configure the time that the switch waits after sending a group-specific query to determine if hosts are still interested in a specific multicast group. The IGMP leave response time can be configured from 100 to 5000 milliseconds. The default leave time is 1000 milliseconds. The timer can be set either globally or on a per-VLAN basis. The VLAN configuration of the leave time overrides the global configuration.

The actual leave latency in the network is usually the configured leave time. However, the leave time *might* vary around the configured time, depending on real-time CPU load conditions, network delays and the amount of traffic sent through the interface.

Note The IGMP configurable leave time is only supported on hosts running IGMP Version 2.

IGMP Report Suppression

Note IGMP report suppression is supported only when the multicast query has IGMPv1 and IGMPv2 reports. This feature is not supported when the query includes IGMPv3 reports.

The switch uses IGMP report suppression to forward only one IGMP report per multicast router query to multicast devices. When IGMP router suppression is enabled (the default), the switch sends the first IGMP report from all hosts for a group to all the multicast routers. The switch does not send the remaining IGMP reports for the group to the multicast routers. This feature prevents duplicate reports from being sent to the multicast devices.

If the multicast router query includes requests only for IGMPv1 and IGMPv2 reports, the switch forwards only the first IGMPv1 or IGMPv2 report from all hosts for a group to all the multicast routers.

If the multicast router query also includes requests for IGMPv3 reports, the switch forwards all IGMPv1, IGMPv2, and IGMPv3 reports for a group to the multicast devices.

If you disable IGMP report suppression, all IGMP reports are forwarded to the multicast routers. For configuration steps, see the "Disabling IGMP Report Suppression" section on page 28-16.

Default IGMP Snooping Configuration

Table 28-3 shows the default IGMP snooping configuration.

Table 28-3 Default IGMP Snooping Configuration

Feature	Default Setting
IGMP snooping	Enabled globally and per VLAN
Multicast routers	None configured
Multicast router learning (snooping) method	PIM-DVMRP
IGMP snooping Immediate Leave	Disabled
Static groups	None configured
TCN[1] flood query count	2
TCN query solicitation	Disabled
IGMP snooping querier	Disabled
IGMP report suppression	Enabled

1. TCN = Topology Change Notification

Snooping Methods

Multicast-capable router ports are added to the forwarding table for every Layer 2 multicast entry. The switch learns of such ports through one of these methods:

- Snooping on IGMP queries, Protocol Independent Multicast (PIM) packets, and Distance Vector Multicast Routing Protocol (DVMRP) packets
- Listening to Cisco Group Management Protocol (CGMP) packets from other routers
- Statically connecting to a multicast router port with the **ip igmp snooping mrouter** global configuration command

You can configure the switch either to snoop on IGMP queries and PIM/DVMRP packets or to listen to CGMP self-join or proxy-join packets. By default, the switch snoops on PIM/DVMRP packets on all VLANs. To learn of multicast router ports through only CGMP packets, use the **ip igmp snooping vlan** *vlan-id* **mrouter learn cgmp** global configuration command. When this command is entered, the router listens to only CGMP self-join and CGMP proxy-join packets and to no other CGMP packets. To learn of multicast router ports through only PIM-DVMRP packets, use the **ip igmp snooping vlan** *vlan-id* **mrouter learn pim-dvmrp** global configuration command.

Note If you want to use CGMP as the learning method and no multicast routers in the VLAN are CGMP proxy-enabled, you must enter the **ip cgmp router-only** command to dynamically access the router.

Multicast Flooding Time After a TCN Event

You can control the time that multicast traffic is flooded after a topology change notification (TCN) event by using the **ip igmp snooping tcn flood query count** global configuration command. This command configures the number of general queries for which multicast data traffic is flooded after a TCN event. Some examples of TCN events are when the client changed its location and the receiver is on same port that was blocked but is now forwarding, and when a port went down without sending a leave message.

If you set the TCN flood query count to 1 by using the **ip igmp snooping tcn flood query count** command, the flooding stops after receiving 1 general query. If you set the count to 7, the flooding continues until 7 general queries are received. Groups are relearned based on the general queries received during the TCN event.

Flood Mode for TCN

When a topology change occurs, the spanning-tree root sends a special IGMP leave message (also known as global leave) with the group multicast address 0.0.0.0. However, when you enable the **ip igmp snooping tcn query solicit** global configuration command, the switch sends the global leave message whether or not it is the spanning-tree root. When the router receives this special leave, it immediately sends general queries, which expedite the process of recovering from the flood mode during the TCN event. Leaves are always sent if the switch is the spanning-tree root regardless of this configuration command. By default, query solicitation is disabled.

Multicast Flooding During a TCN Event

When the switch receives a TCN, multicast traffic is flooded to all the ports until 2 general queries are received. If the switch has many ports with attached hosts that are subscribed to different multicast groups, this flooding might exceed the capacity of the link and cause packet loss. You can use the **ip igmp snooping tcn flood** interface configuration command to control this behavior.

IGMP Snooping Querier Guidelines

- Configure the VLAN in global configuration mode.
- Configure an IP address on the VLAN interface. When enabled, the IGMP snooping querier uses the IP address as the query source address.
- If there is no IP address configured on the VLAN interface, the IGMP snooping querier tries to use the configured global IP address for the IGMP querier. If there is no global IP address specified, the IGMP querier tries to use the VLAN switch virtual interface (SVI) IP address (if one exists). If there

is no SVI IP address, the switch uses the first available IP address configured on the switch. The first IP address available appears in the output of the **show ip interface** privileged EXEC command. The IGMP snooping querier does not generate an IGMP general query if it cannot find an available IP address on the switch.

- The IGMP snooping querier supports IGMP Versions 1 and 2.
- When administratively enabled, the IGMP snooping querier moves to the nonquerier state if it detects the presence of a multicast router in the network.
- When it is administratively enabled, the IGMP snooping querier moves to the operationally disabled state under these conditions:
  - IGMP snooping is disabled in the VLAN.
  - PIM is enabled on the SVI of the corresponding VLAN.

IGMP Report Suppression

IGMP report suppression is enabled by default. When it is enabled, the switch forwards only one IGMP report per multicast router query. When report suppression is disabled, all IGMP reports are forwarded to the multicast routers.

Multicast VLAN Registration

Note To use this feature, the switch must be running the LAN Base image.

Multicast VLAN Registration (MVR) is designed for applications using wide-scale deployment of multicast traffic across an Ethernet ring-based service-provider network (for example, the broadcast of multiple television channels over a service-provider network). MVR allows a subscriber on a port to subscribe and unsubscribe to a multicast stream on the network-wide multicast VLAN. It allows the single multicast VLAN to be shared in the network while subscribers remain in separate VLANs. MVR provides the ability to continuously send multicast streams in the multicast VLAN, but to isolate the streams from the subscriber VLANs for bandwidth and security reasons.

MVR assumes that subscriber ports subscribe and unsubscribe (join and leave) these multicast streams by sending out IGMP join and leave messages. These messages can originate from an IGMP Version-2-compatible host with an Ethernet connection. Although MVR operates on the underlying mechanism of IGMP snooping, the two features operate independently of each other. One can be enabled or disabled without affecting the behavior of the other feature. However, if IGMP snooping and MVR are both enabled, MVR reacts only to join and leave messages from multicast groups configured under MVR. Join and leave messages from all other multicast groups are managed by IGMP snooping.

The switch CPU identifies the MVR IP multicast streams and their associated IP multicast group in the switch forwarding table, intercepts the IGMP messages, and modifies the forwarding table to include or remove the subscriber as a receiver of the multicast stream, even though the receivers might be in a different VLAN from the source. This forwarding behavior selectively allows traffic to cross between different VLANs.

You can set the switch for compatible or dynamic mode of MVR operation:

- In compatible mode, multicast data received by MVR hosts is forwarded to all MVR data ports, regardless of MVR host membership on those ports. The multicast data is forwarded only to those receiver ports that MVR hosts have joined, either by IGMP reports or by MVR static configuration. IGMP reports received from MVR hosts are never forwarded from MVR data ports that were configured in the switch.
- In dynamic mode, multicast data received by MVR hosts on the switch is forwarded from only those MVR data and client ports that the MVR hosts have joined, either by IGMP reports or by MVR static configuration. Any IGMP reports received from MVR hosts are also forwarded from all the MVR data ports in the switch. This eliminates using unnecessary bandwidth on MVR data port links, which occurs when the switch runs in compatible mode.

Only Layer 2 ports take part in MVR. You must configure ports as MVR receiver ports. Only one MVR multicast VLAN per switch is supported.

MVR in a Multicast Television Application

In a multicast television application, a PC or a television with a set-top box can receive the multicast stream. Multiple set-top boxes or PCs can be connected to one subscriber port, which is a switch port configured as an MVR receiver port. Figure 28-3 is an example configuration. DHCP assigns an IP address to the set-top box or the PC. When a subscriber selects a channel, the set-top box or PC sends an IGMP report to Switch A to join the appropriate multicast. If the IGMP report matches one of the configured IP multicast group addresses, the switch CPU modifies the hardware address table to include this receiver port and VLAN as a forwarding destination of the specified multicast stream when it is received from the multicast VLAN. Uplink ports that send and receive multicast data to and from the multicast VLAN are called MVR source ports.

Figure 28-3 Multicast VLAN Registration Example

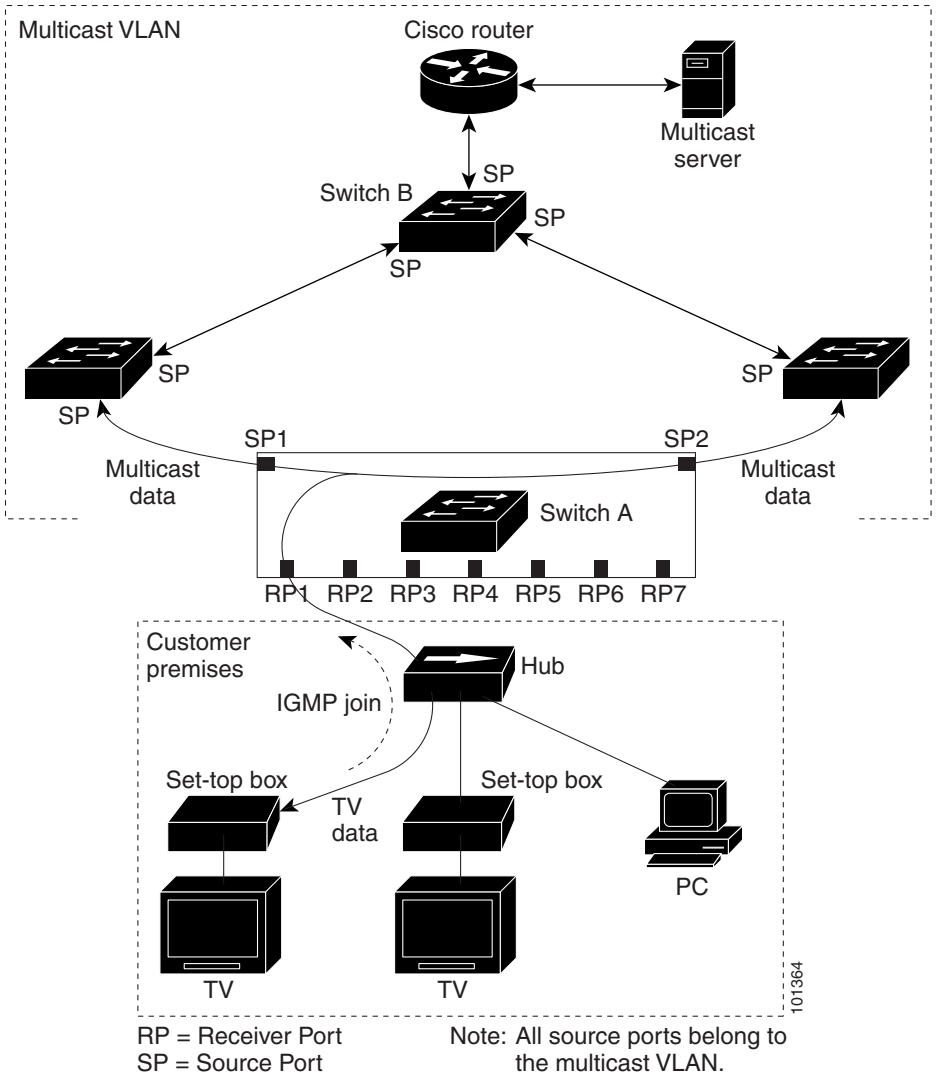

When a subscriber changes channels or turns off the television, the set-top box sends an IGMP leave message for the multicast stream. The switch CPU sends a MAC-based general query through the receiver port VLAN. If there is another set-top box in the VLAN still subscribing to this group, that set-top box must respond within the maximum response time specified in the query. If the CPU does not receive a response, it eliminates the receiver port as a forwarding destination for this group.

Without Immediate Leave, when the switch receives an IGMP leave message from a subscriber on a receiver port, it sends out an IGMP query on that port and waits for IGMP group membership reports. If no reports are received in a configured time period, the receiver port is removed from multicast group membership. With Immediate Leave, an IGMP query is not sent from the receiver port on which the IGMP leave was received. As soon as the leave message is received, the receiver port is removed from multicast group membership, which speeds up leave latency. Enable the Immediate Leave feature only on receiver ports to which a single receiver device is connected.

MVR eliminates the need to duplicate television-channel multicast traffic for subscribers in each VLAN. Multicast traffic for all channels is only sent around the VLAN trunk once—only on the multicast VLAN. The IGMP leave and join messages are in the VLAN to which the subscriber port is assigned.

These messages dynamically register for streams of multicast traffic in the multicast VLAN on the Layer 3 device. Switch B. The access layer switch, Switch A, modifies the forwarding behavior to allow the traffic to be forwarded from the multicast VLAN to the subscriber port in a different VLAN, selectively allowing traffic to cross between two VLANs.

IGMP reports are sent to the same IP multicast group address as the multicast data. The Switch A CPU must capture all IGMP join and leave messages from receiver ports and forward them to the multicast VLAN of the source (uplink) port, based on the MVR mode.

Default MVR Settings

Table 28-4 Default MVR Settings

Feature	Default Setting
MVR	Disabled globally and per interface
Multicast addresses	None configured
Query response time	0.5 second
Multicast VLAN	VLAN 1
Mode	Compatible
Interface (per port) default	Neither a receiver nor a source port
Immediate Leave	Disabled on all ports

MVR Configuration Guidelines and Limitations

- Receiver ports can only be access ports; they cannot be trunk ports. Receiver ports on a switch can be in different VLANs, but should not belong to the multicast VLAN.
- The maximum number of multicast entries (MVR group addresses) that can be configured on a switch (that is, the maximum number of television channels that can be received) is 256.
- MVR multicast data received in the source VLAN and leaving from receiver ports has its time-to-live (TTL) decremented by 1 in the switch.
- Because MVR on the switch uses IP multicast addresses instead of MAC multicast addresses, aliased IP multicast addresses are allowed on the switch. However, if the switch is interoperating with Catalyst 3550 or Catalyst 3500 XL switches, you should not configure IP addresses that alias between themselves or with the reserved IP multicast addresses (in the range 224.0.0.xxx).
- Do not configure MVR on private VLAN ports.
- MVR is not supported when multicast routing is enabled on a switch. If you enable multicast routing and a multicast routing protocol while MVR is enabled, MVR is disabled, and you receive a warning message. If you try to enable MVR while multicast routing and a multicast routing protocol are enabled, the operation to enable MVR is cancelled, and you receive an error message.
- MVR can coexist with IGMP snooping on a switch.
- MVR data received on an MVR receiver port is not forwarded to MVR source ports.
- MVR does not support IGMPv3 messages.

IGMP Filtering and Throttling

In some environments, for example, metropolitan or multiple-dwelling unit (MDU) installations, you might want to control the set of multicast groups to which a user on a switch port can belong. You can control the distribution of multicast services, such as IP/TV, based on some type of subscription or service plan. You might also want to limit the number of multicast groups to which a user on a switch port can belong.

With the IGMP filtering feature, you can filter multicast joins on a per-port basis by configuring IP multicast profiles and associating them with individual switch ports. An IGMP profile can contain one or more multicast groups and specifies whether access to the group is permitted or denied. If an IGMP profile denying access to a multicast group is applied to a switch port, the IGMP join report requesting the stream of IP multicast traffic is dropped, and the port is not allowed to receive IP multicast traffic from that group. If the filtering action permits access to the multicast group, the IGMP report from the port is forwarded for normal processing. You can also set the maximum number of IGMP groups that a Layer 2 interface can join.

IGMP filtering controls only group-specific query and membership reports, including join and leave reports. It does not control general IGMP queries. IGMP filtering has no relationship with the function that directs the forwarding of IP multicast traffic. The filtering feature operates in the same manner whether CGMP or MVR is used to forward the multicast traffic.

IGMP filtering is applicable only to the dynamic learning of IP multicast group addresses, not static configuration.

With the IGMP throttling feature, you can set the maximum number of IGMP groups that a Layer 2 interface can join. If the maximum number of IGMP groups is set, the IGMP snooping forwarding table contains the maximum number of entries, and the interface receives an IGMP join report, you can configure an interface to drop the IGMP report or to replace the randomly selected multicast entry with the received IGMP report.

Note IGMPv3 join and leave messages are not supported on switches running IGMP filtering.

Default IGMP Filtering and Throttling Configuration

Table 28-5 shows the default IGMP filtering configuration.

Table 28-5 Default IGMP Filtering Configuration

Feature	Default Setting
IGMP filters	None applied
IGMP maximum number of IGMP groups	No maximum set
IGMP profiles	None defined
IGMP profile action	Deny the range addresses

When the maximum number of groups is in forwarding table, the default IGMP throttling action is to deny the IGMP report.

IGMP Profiles

To configure an IGMP profile, use the **ip igmp profile** global configuration command with a profile number to create an IGMP profile and to enter IGMP profile configuration mode. From this mode, you can specify the parameters of the IGMP profile to be used for filtering IGMP join requests from a port. When you are in IGMP profile configuration mode, you can create the profile by using these commands:

- **deny**—Specifies that matching addresses are denied; this is the default.
- **exit**—Exits from igmp-profile configuration mode.
- **no**—Negates a command or returns to its defaults.
- **permit**—Specifies that matching addresses are permitted.
- **range**—Specifies a range of IP addresses for the profile. You can enter a single IP address or a range with a start and an end address.

The default is for the switch to have no IGMP profiles configured. When a profile is configured, if neither the **permit** nor **deny** keyword is included, the default is to deny access to the range of IP addresses.

To control access as defined in an IGMP profile, use the **ip igmp filter** interface configuration command to apply the profile to the appropriate interfaces. You can apply IGMP profiles only to Layer 2 access ports; you cannot apply IGMP profiles to routed ports or SVIs. You cannot apply profiles to ports that belong to an EtherChannel port group. You can apply a profile to multiple interfaces, but each interface can have only one profile applied to it.

IGMP Throttling Action

After you set the maximum number of IGMP groups that a Layer 2 interface can join, you can configure an interface to replace the existing group with the new group for which the IGMP report was received by using the **ip igmp max-groups action replace** interface configuration command. Use the **no** form of this command to return to the default, which is to drop the IGMP join report.

Follow these guidelines when configuring the IGMP throttling action:

- This restriction can be applied only to Layer 2 ports. You can use this command on a logical EtherChannel interface but cannot use it on ports that belong to an EtherChannel port group.
- When the maximum group limitation is set to the default (no maximum), entering the **ip igmp max-groups action** {**deny** | **replace**} command has no effect.
- If you configure the throttling action and set the maximum group limitation after an interface has added multicast entries to the forwarding table, the forwarding-table entries are either aged out or removed, depending on the throttling action.
  - If you configure the throttling action as **deny**, the entries that were previously in the forwarding table are not removed but are aged out. After these entries are aged out and the maximum number of entries is in the forwarding table, the switch drops the next IGMP report received on the interface.
  - If you configure the throttling action as **replace**, the entries that were previously in the forwarding table are removed. When the maximum number of entries is in the forwarding table, the switch replaces a randomly selected entry with the received IGMP report.

  To prevent the switch from removing the forwarding-table entries, you can configure the IGMP throttling action before an interface adds entries to the forwarding table.

How to Configure IGMP Snooping and MVR

Configuring IGMP Snooping

Enabling or Disabling IGMP Snooping

By default, IGMP snooping is globally enabled on the switch. When globally enabled or disabled, it is also enabled or disabled in all existing VLAN interfaces. IGMP snooping is by default enabled on all VLANs, but can be enabled and disabled on a per-VLAN basis.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip igmp snooping**	Globally enables IGMP snooping in all existing VLAN interfaces.
	or	or
	ip igmp snooping vlan *vlan-id*	Enables IGMP snooping on the VLAN interface. The VLAN ID range is 1 to 1001 and 1006 to 4096. IGMP snooping must be globally enabled before you can enable VLAN snooping.
Step 3	**end**	Returns to privileged EXEC mode.

Setting IGMP Snooping Parameters

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip igmp snooping vlan** *vlan-id* **mrouter learn** {**cgmp** | **pim-dvmrp**}	(Optional) Enables IGMP snooping on a VLAN. The VLAN ID range is 1 to 1001 and 1006 to 4096. Specifies the multicast router learning method: • **cgmp**—Listens for CGMP packets. This method is useful for reducing control traffic. • **pim-dvmrp**—Snoops on IGMP queries and PIM-DVMRP packets. This is the default.
Step 3	**ip igmp snooping vlan** *vlan-id* **mrouter interface** *interface-id*	Adds a multicast router port (adds a static connection to a multicast router). (Optional) Specifies the multicast router VLAN ID and the interface to the multicast router. • The VLAN ID range is 1 to 1001 and 1006 to 4096. • The interface can be a physical interface or a port channel. The port-channel range is 1 to 6. • Static connections to multicast routers are supported only on switch ports.

	Command	Purpose
Step 4	**ip igmp snooping vlan** *vlan-id* **static** *ip_address* **interface** *interface-id*	(Optional) Statically configures a Layer 2 port as a member of a multicast group: • *vlan-id*—Multicast group VLAN ID. The range is 1 to 1001 and 1006 to 4096. • *ip-address*—Group IP address. • *interface-id*—Member port. It can be a physical interface or a port channel (1 to 6).
Step 5	**ip igmp snooping vlan** *vlan-id* **immediate-leave**	(Optional) Enables IGMP Immediate Leave on the VLAN interface. **Note** Immediate Leave is supported only on IGMP Version 2 hosts.
Step 6	**ip igmp snooping last-member-query-interval** *time*	(Optional) Configures the IGMP leave timer globally. The range is 100 to 32768 milliseconds. The default is 1000 seconds.
Step 7	**ip igmp snooping vlan** *vlan-id* **last-member-query-interval** *time*	(Optional) Configures the IGMP leave time on the VLAN interface. The range is 100 to 32768 milliseconds. **Note** Configuring the leave time on a VLAN overrides the globally configured timer.
Step 8	**end**	Returns to privileged EXEC mode.

Configuring TCN

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip igmp snooping tcn flood query count** *count*	Specifies the number of IGMP general queries for which the multicast traffic is flooded. The range is 1 to 10. By default, the flooding query count is 2.
Step 3	**ip igmp snooping tcn query solicit**	Sends an IGMP leave message (global leave) to speed the process of recovering from the flood mode caused during a TCN event. By default, query solicitation is disabled. **Note** Enable the switch to send the global leave message whether or not it is the spanning-tree root.
Step 4	**interface** *interface-id*	Specifies the interface to be configured, and enter interface configuration mode.
Step 5	**no ip igmp snooping tcn flood**	Disables the flooding of multicast traffic during a spanning-tree TCN event. By default, multicast flooding is enabled on an interface.
Step 6	**end**	Returns to privileged EXEC mode.

Configuring the IGMP Snooping Querier

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip igmp snooping querier**	Enables the IGMP snooping querier.
Step 3	**ip igmp snooping querier address** *ip_address*	(Optional) Specifies an IP address for the IGMP snooping querier. If you do not specify an IP address, the querier tries to use the global IP address configured for the IGMP querier. **Note** The IGMP snooping querier does not generate an IGMP general query if it cannot find an IP address on the switch.
Step 4	**ip igmp snooping querier query-interval** *interval-count*	(Optional) Sets the interval between IGMP queriers. The range is 1 to 18000 seconds.
Step 5	**ip igmp snooping querier tcn query** [**count** *count* | **interval** *interval*]	(Optional) Sets the time between Topology Change Notification (TCN) queries. The count range is 1 to 10. The interval range is 1 to 255 seconds.
Step 6	**ip igmp snooping querier timer expiry** *timeout*	(Optional) Sets the length of time until the IGMP querier expires. The range is 60 to 300 seconds.
Step 7	**ip igmp snooping querier version** *version*	(Optional) Selects the IGMP version number that the querier feature uses. Select 1 or 2.
Step 8	**end**	Returns to privileged EXEC mode.

Disabling IGMP Report Suppression

Before You Begin

IGMP report suppression is supported only when the multicast query has IGMPv1 and IGMPv2 reports. This feature is not supported when the query includes IGMPv3 reports.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**no ip igmp snooping report-suppression**	Disables IGMP report suppression.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring MVR

Configuring MVR Global Parameters

You do not need to set the optional MVR parameters if you choose to use the default settings. If you do want to change the default parameters (except for the MVR VLAN), you must first enable MVR.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mvr**	Enables MVR on the switch.
Step 3	**mvr group** *ip-address* [*count*]	Configures an IP multicast address on the switch or use the *count* parameter to configure a contiguous series of MVR group addresses (the range for *count* is 1 to 256; the default is 1). Any multicast data sent to this address is sent to all source ports on the switch and all receiver ports that have elected to receive data on that multicast address. Each multicast address would correspond to one television channel.
Step 4	**mvr querytime** *value*	(Optional) Defines the maximum time to wait for IGMP report memberships on a receiver port before removing the port from multicast group membership. The value is in units of tenths of a second. The range is 1 to 100, and the default is 5 tenths or one-half second.
Step 5	**mvr vlan** *vlan-id*	(Optional) Specifies the VLAN in which multicast data is received; all source ports must belong to this VLAN. The VLAN range is 1 to 1001 and 1006 to 4096. The default is VLAN 1.
Step 6	**mvr mode** {**dynamic** | **compatible**}	(Optional) Specifies the MVR mode of operation: • **dynamic**—Allows dynamic MVR membership on source ports. • **compatible**—Is compatible with Catalyst 3500 XL and Catalyst 2900 XL switches and does not support IGMP dynamic joins on source ports. The default is **compatible** mode.
Step 7	**end**	Returns to privileged EXEC mode.

Configuring MVR Interfaces

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mvr**	Enables MVR on the switch.
Step 3	**interface** *interface-id*	Specifies the Layer 2 port to configure, and enters interface configuration mode.
Step 4	**mvr type** {**source** | **receiver**}	Configures an MVR port as one of these: • **source**—Configures uplink ports that receive and send multicast data as source ports. Subscribers cannot be directly connected to source ports. All source ports on a switch belong to the single multicast VLAN. • **receiver**—Configures a port as a receiver port if it is a subscriber port and should only receive multicast data. It does not receive data unless it becomes a member of the multicast group, either statically or by using IGMP leave and join messages. Receiver ports cannot belong to the multicast VLAN. The default configuration is as a non-MVR port. If you attempt to configure a non-MVR port with MVR characteristics, the operation fails.

	Command	Purpose
Step 5	**mvr vlan** *vlan-id* **group** [*ip-address*]	(Optional) Statically configures a port to receive multicast traffic sent to the multicast VLAN and the IP multicast address. A port statically configured as a member of a group remains a member of the group until statically removed. **Note** In compatible mode, this command applies to only receiver ports. In dynamic mode, it applies to receiver ports and source ports. Receiver ports can also dynamically join multicast groups by using IGMP join and leave messages.
Step 6	**mvr immediate**	(Optional) Enables the Immediate-Leave feature of MVR on the port. **Note** This command applies to only receiver ports and should only be enabled on receiver ports to which a single receiver device is connected.
Step 7	**end**	Returns to privileged EXEC mode.

Configuring IGMP

Configuring IGMP Profiles

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip igmp profile** *profile number*	Assigns a number to the profile you are configuring, and enter IGMP profile configuration mode. The profile number range is 1 to 4294967295.
Step 3	**permit** | **deny**	(Optional) Sets the action to permit or deny access to the IP multicast address. If no action is configured, the default for the profile is to deny access.
Step 4	**range** *ip multicast address*	Enters the IP multicast address or range of IP multicast addresses to which access is being controlled. If entering a range, enter the low IP multicast address, a space, and the high IP multicast address. You can use the **range** command multiple times to enter multiple addresses or ranges of addresses.
Step 5	**end**	Returns to privileged EXEC mode.

Configuring IGMP Interfaces

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the physical interface, and enter interface configuration mode. The interface must be a Layer 2 port that does not belong to an EtherChannel port group.

	Command	Purpose
Step 3	**ip igmp filter** *profile number*	Applies the specified IGMP profile to the interface. The range is 1 to 4294967295.
Step 4	**ip igmp max-groups** *number*	Sets the maximum number of IGMP groups that the interface can join. The range is 0 to 4294967294. The default is to have no maximum set.
Step 5	**ip igmp max-groups action** {**deny** | **replace**}	When an interface receives an IGMP report and the maximum number of entries is in the forwarding table, specify the action that the interface takes: • **deny**—Drops the report. • **replace**—Replaces the existing group with the new group for which the IGMP report was received.
Step 6	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining IGMP Snooping and MVR

Command	Purpose
show ip igmp snooping [**vlan** *vlan-id*]	Displays the snooping configuration information for all VLANs on the switch or for a specified VLAN. (Optional) Enter **vlan** *vlan-id* to display information for a single VLAN. The VLAN ID range is 1 to 1001 and 1006 to 4096.
show ip igmp snooping groups [**count** |**dynamic** [**count**] | **user** [**count**]]	Displays multicast table information for the switch or about a specific parameter: • **count**—Displays the total number of entries for the specified command options instead of the actual entries. • **dynamic**—Displays entries learned through IGMP snooping. • **user**—Displays only the user-configured multicast entries.
show ip igmp snooping groups vlan *vlan-id* [*ip_address* | **count** | **dynamic** [**count**] | **user**[**count**]]	Displays multicast table information for a multicast VLAN or about a specific parameter for the VLAN: • *vlan-id*—The VLAN ID range is 1 to 1001 and 1006 to 4096. • **count**—Displays the total number of entries for the specified command options instead of the actual entries. • **dynamic**—Displays entries learned through IGMP snooping. • *ip_address*—Displays characteristics of the multicast group with the specified group IP address. • **user**—Displays only the user-configured multicast entries.

Command	Purpose
show ip igmp snooping mrouter [**vlan** *vlan-id*]	Displays information on dynamically learned and manually configured multicast router interfaces. **Note** When you enable IGMP snooping, the switch automatically learns the interface to which a multicast router is connected. These are dynamically learned interfaces. (Optional) Enter **vlan** *vlan-id* to display information for a single VLAN.
show ip igmp snooping querier [**vlan** *vlan-id*]	Displays information about the IP address and receiving port for the most-recently received IGMP query messages in the VLAN. (Optional) Enter **vlan** *vlan-id* to display information for a single VLAN.
show ip igmp snooping querier [**vlan** *vlan-id*] **detail**	Displays information about the IP address and receiving port of the most-recently received IGMP query message in the VLAN and the configuration and operational state of the IGMP snooping querier in the VLAN.
show ip igmp profile [*profile number*]	Displays the specified IGMP profile or all the IGMP profiles defined on the switch.
show mvr	Displays MVR status and values for the switch—whether MVR is enabled or disabled, the multicast VLAN, the maximum (256) and current (0 through 256) number of multicast groups, the query response time, and the MVR mode.
show mvr interface [*interface-id*] [**members** [**vlan** *vlan-**id***]]	Displays all MVR interfaces and their MVR configurations. When a specific interface is entered, displays this information: • Type—Receiver or Source • Status—One of these: – Active means the port is part of a VLAN. – Up/Down means that the port is forwarding or nonforwarding. – Inactive means that the port is not part of any VLAN. • Immediate Leave—Enabled or Disabled If the **members** keyword is entered, displays all multicast group members on this port or, if a VLAN identification is entered, all multicast group members on the VLAN. The VLAN ID range is 1 to 1001 and 1006 to 4096.
show mvr members [*ip-address*]	Displays all receiver and source ports that are members of any IP multicast group or the specified IP multicast group IP address.
show ip igmp profile *profile number*	Verifies the profile configuration.
show ip igmp snooping mrouter [**vlan** *vlan-id*]	Verifies that IGMP snooping is enabled on the VLAN interface.

Configuration Examples for IGMP Snooping

Configuring IGMP Snooping: Example

This example shows how to configure IGMP snooping to use CGMP packets as the learning method:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 1 mrouter learn cgmp
Switch(config)# end
```

Disabling a Multicast Router Port: Example

To remove a multicast router port from the VLAN, use the **no ip igmp snooping vlan** *vlan-id* **mrouter interface** *interface-id* global configuration command.

This example shows how to enable a static connection to a multicast router:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 200 mrouter interface gigabitethernet1/2
Switch(config)# end
```

Statically Configuring a Host on a Port: Example

This example shows how to statically configure a host on a port:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 105 static 224.2.4.12 interface gigabitethernet1/1
Switch(config)# end
```

Enabling IGMP Immediate Leave: Example

This example shows how to enable IGMP Immediate Leave on VLAN 130:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 130 immediate-leave
Switch(config)# end
```

Setting the IGMP Snoopng Querier Parameters: Examples

This example shows how to set the IGMP snooping querier source address to 10.0.0.64:

```
Switch# configure terminal
Switch(config)# ip igmp snooping querier 10.0.0.64
Switch(config)# end
```

This example shows how to set the IGMP snooping querier maximum response time to 25 seconds:

```
Switch# configure terminal
Switch(config)# ip igmp snooping querier query-interval 25
Switch(config)# end
```

This example shows how to set the IGMP snooping querier timeout to 60 seconds:

```
Switch# configure terminal
```

```
Switch(config)# ip igmp snooping querier timeout expiry 60
Switch(config)# end
```

This example shows how to set the IGMP snooping querier feature to version 2:

```
Switch# configure terminal
Switch(config)# no ip igmp snooping querier version 2
Switch(config)# end
```

Enabling MVR: Examples

This example shows how to enable MVR, configure the group address, set the query time to 1 second (10 tenths), specify the MVR multicast VLAN as VLAN 22, and set the MVR mode as dynamic:

```
Switch(config)# mvr
Switch(config)# mvr group 228.1.23.4
Switch(config)# mvr querytime 10
Switch(config)# mvr vlan 22
Switch(config)# mvr mode dynamic
Switch(config)# end
```

You can use the **show mvr members** privileged EXEC command to verify the MVR multicast group addresses on the switch.

This example shows how to configure a port as a receiver port, statically configure the port to receive multicast traffic sent to the multicast group address, configure Immediate Leave on the port, and verify the results:

```
Switch(config)# mvr
Switch(config)# interface gigabitethernet1/2
Switch(config-if)# mvr type receiver
Switch(config-if)# mvr vlan 22 group 228.1.23.4
Switch(config-if)# mvr immediate
Switch(config)# end
Switch# show mvr interface
Port    Type        Status          Immediate Leave
----    ----        -------         ---------------
Gi1/2   RECEIVER    ACTIVE/DOWN     ENABLED
```

Creating an IGMP Profile: Example

This example shows how to create IGMP profile 4 allowing access to the single IP multicast address and how to verify the configuration. If the action was to deny (the default), it would not appear in the **show ip igmp profile** output display.

```
Switch(config)# ip igmp profile 4
Switch(config-igmp-profile)# permit
Switch(config-igmp-profile)# range 229.9.9.0
Switch(config-igmp-profile)# end
Switch# show ip igmp profile 4
IGMP Profile 4
    permit
    range 229.9.9.0 229.9.9.0
```

Applying an IGMP Profile: Example

This example shows how to apply IGMP profile 4 to a port:

```
Switch(config)# interface gigabitethernet1/2
Switch(config-if)# ip igmp filter 4
Switch(config-if)# end
```

Limiting IGMP Groups: Example

This example shows how to limit to 25 the number of IGMP groups that a port can join:

```
Switch(config)# interface gigabitethernet1/2
Switch(config-if)# ip igmp max-groups 25
Switch(config-if)# end
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
Cisco IOS multicast commands	*Cisco IOS IP Command Reference, Volume 3 of 3:Multicast*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 29

Configuring Port-Based Traffic Control

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Restrictions for Port-Based Traffic Control

- To use this feature, the switch must be running the LAN Base image.

Information About Port-Based Traffic Control

Storm Control

Storm control prevents traffic on a LAN from being disrupted by a broadcast, multicast, or unicast storm on one of the physical interfaces. A LAN storm occurs when packets flood the LAN, creating excessive traffic and degrading network performance. Errors in the protocol-stack implementation, mistakes in network configurations, or users issuing a denial-of-service attack can cause a storm.

Storm control (or traffic suppression) monitors packets passing from an interface to the switching bus and determines if the packet is unicast, multicast, or broadcast. The switch counts the number of packets of a specified type received within the 1-second time interval and compares the measurement with a predefined suppression-level threshold.

Storm control uses one of these methods to measure traffic activity:

- Bandwidth as a percentage of the total available bandwidth of the port that can be used by the broadcast, multicast, or unicast traffic
- Traffic rate in packets per second at which broadcast, multicast, or unicast packets are received.
- Traffic rate in bits per second at which broadcast, multicast, or unicast packets are received.
- Traffic rate in packets per second and for small frames. This feature is enabled globally. The threshold for small frames is configured for each interface.

With each method, the port blocks traffic when the rising threshold is reached. The port remains blocked until the traffic rate drops below the falling threshold (if one is specified) and then resumes normal forwarding. If the falling suppression level is not specified, the switch blocks all traffic until the traffic rate drops below the rising suppression level. In general, the higher the level, the less effective the protection against broadcast storms.

Note When the storm control threshold for multicast traffic is reached, all multicast traffic except control traffic, such as bridge protocol data unit (BDPU) and Cisco Discovery Protocol (CDP) frames, are blocked. However, the switch does not differentiate between routing updates, such as OSPF, and regular multicast data traffic, so both types of traffic are blocked.

The graph in Figure 29-1 shows broadcast traffic patterns on an interface over a given period of time. The example can also be applied to multicast and unicast traffic. In this example, the broadcast traffic being forwarded exceeded the configured threshold between time intervals T1 and T2 and between T4 and T5. When the amount of specified traffic exceeds the threshold, all traffic of that kind is dropped for the next time period. Therefore, broadcast traffic is blocked during the intervals following T2 and T5. At the next time interval (for example, T3), if broadcast traffic does not exceed the threshold, it is again forwarded.

Figure 29-1 Broadcast Storm Control Example

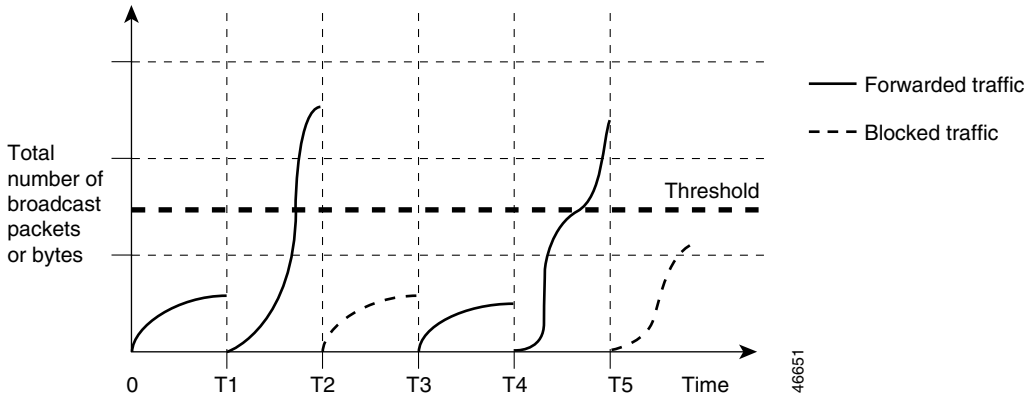

The combination of the storm-control suppression level and the 1-second time interval controls the way the storm control algorithm works. A higher threshold allows more packets to pass through. A threshold value of 100 percent means that no limit is placed on the traffic. A value of 0.0 means that all broadcast, multicast, or unicast traffic on that port is blocked.

Note Because packets do not arrive at uniform intervals, the 1-second time interval during which traffic activity is measured can affect the behavior of storm control.

You use the **storm-control** interface configuration commands to set the threshold value for each traffic type.

Default Storm Control Configuration

By default, unicast, broadcast, and multicast storm control are disabled on the switch interfaces; that is, the suppression level is 100 percent.

Storm Control and Threshold Levels

You configure storm control on a port and enter the threshold level that you want to be used for a particular type of traffic.

However, because of hardware limitations and the way in which packets of different sizes are counted, threshold percentages are approximations. Depending on the sizes of the packets making up the incoming traffic, the actual enforced threshold might differ from the configured level by several percentage points.

Note Storm control is supported on physical interfaces. You can also configure storm control on an EtherChannel. When storm control is configured on an EtherChannel, the storm control settings propagate to the EtherChannel physical interfaces.

Small-Frame Arrival Rate

Incoming VLAN-tagged packets smaller than 67 bytes are considered *small frames*. They are forwarded by the switch, but they do not cause the switch storm-control counters to increment. In Cisco IOS Release 12.2(44)SE and later, you can configure a port to be error disabled if small frames arrive at a specified rate (threshold).

You globally enable the small-frame arrival feature on the switch and then configure the small-frame threshold for packets on each interface. Packets smaller than the minimum size and arriving at a specified rate (the threshold) are dropped since the port is error disabled.

If the **errdisable recovery cause small-frame** global configuration command is entered, the port is reenabled after a specified time. (You specify the recovery time by using **errdisable recovery** global configuration command.)

Protected Ports

Some applications require that no traffic be forwarded at Layer 2 between ports on the same switch so that one neighbor does not see the traffic generated by another neighbor. In such an environment, the use of protected ports ensures that there is no exchange of unicast, broadcast, or multicast traffic between these ports on the switch.

Protected ports have these features:

- A protected port does not forward any traffic (unicast, multicast, or broadcast) to any other port that is also a protected port. Data traffic cannot be forwarded between protected ports at Layer 2; only control traffic, such as PIM packets, is forwarded because these packets are processed by the CPU and forwarded in software. All data traffic passing between protected ports must be forwarded through a Layer 3 device.
- Forwarding behavior between a protected port and a nonprotected port proceeds as usual.

Protected Port Configuration Guidelines

You can configure protected ports on a physical interface (for example, Gigabit Ethernet port 1) or an EtherChannel group (for example, port-channel 5). When you enable protected ports for a port channel, it is enabled for all ports in the port-channel group.

Do not configure a private-VLAN port as a protected port. Do not configure a protected port as a private-VLAN port. A private-VLAN isolated port does not forward traffic to other isolated ports or community ports. For more information about private VLANs, see Chapter 19, "Configuring Private VLANs."

Port Blocking

By default, the switch floods packets with unknown destination MAC addresses out of all ports. If unknown unicast and multicast traffic is forwarded to a protected port, there could be security issues. To prevent unknown unicast or multicast traffic from being forwarded from one port to another, you can block a port (protected or nonprotected) from flooding unknown unicast or multicast packets to other ports.

Note With multicast traffic, the port blocking feature blocks only pure Layer 2 packets. Multicast packets that contain IPv4 or IPv6 information in the header are not blocked.

Port Security

You can use the port security feature to restrict input to an interface by limiting and identifying MAC addresses of the stations allowed to access the port. When you assign secure MAC addresses to a secure port, the port does not forward packets with source addresses outside the group of defined addresses. If you limit the number of secure MAC addresses to one and assign a single secure MAC address, the workstation attached to that port is assured the full bandwidth of the port.

If a port is configured as a secure port and the maximum number of secure MAC addresses is reached, when the MAC address of a station attempting to access the port is different from any of the identified secure MAC addresses, a security violation occurs. Also, if a station with a secure MAC address configured or learned on one secure port attempts to access another secure port, a violation is flagged.

Secure MAC Addresses

You configure the maximum number of secure addresses allowed on a port by using the **switchport port-security maximum** *value* interface configuration command.

Note If you try to set the maximum value to a number less than the number of secure addresses already configured on an interface, the command is rejected.

The switch supports these types of secure MAC addresses:

- Static secure MAC addresses—These are manually configured by using the **switchport port-security mac-address** *mac-address* interface configuration command, stored in the address table, and added to the switch running configuration.
- Dynamic secure MAC addresses—These are dynamically configured, stored only in the address table, and removed when the switch restarts.
- *Sticky* secure MAC addresses—These can be dynamically learned or manually configured, stored in the address table, and added to the running configuration. If these addresses are saved in the configuration file, when the switch restarts, the interface does not need to dynamically reconfigure them.

You can configure an interface to convert the dynamic MAC addresses to sticky secure MAC addresses and to add them to the running configuration by enabling *sticky learning*. To enable sticky learning, enter the **switchport port-security mac-address sticky** interface configuration command. When you enter this command, the interface converts all the dynamic secure MAC addresses, including those that were dynamically learned before sticky learning was enabled, to sticky secure MAC addresses. All sticky secure MAC addresses are added to the running configuration.

The sticky secure MAC addresses do not automatically become part of the configuration file, which is the startup configuration used each time the switch restarts. If you save the sticky secure MAC addresses in the configuration file, when the switch restarts, the interface does not need to relearn these addresses. If you do not save the sticky secure addresses, they are lost.

If sticky learning is disabled, the sticky secure MAC addresses are converted to dynamic secure addresses and are removed from the running configuration.

The maximum number of secure MAC addresses that you can configure on a switch is set by the maximum number of available MAC addresses allowed in the system. This number is determined by the active Switch Database Management (SDM) template. See Chapter 11, "Configuring SDM Templates." This number is the total of available MAC addresses, including those used for other Layer 2 functions and any other secure MAC addresses configured on interfaces.

Security Violations

It is a security violation when one of these situations occurs:

- The maximum number of secure MAC addresses have been added to the address table, and a station whose MAC address is not in the address table attempts to access the interface.
- An address learned or configured on one secure interface is seen on another secure interface in the same VLAN and on the same switch.

You can configure the interface for one of four violation modes, based on the action to be taken if a violation occurs:

- protect—When the number of secure MAC addresses reaches the maximum limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses to drop below the maximum value or increase the number of maximum allowable addresses. You are not notified that a security violation has occurred.

  Note We do not recommend configuring the protect violation mode on a trunk port. The protect mode disables learning when any VLAN reaches its maximum limit, even if the port has not reached its maximum limit.

- restrict—When the number of secure MAC addresses reaches the maximum limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses to drop below the maximum value or increase the number of maximum allowable addresses. In this mode, you are notified that a security violation has occurred. An SNMP trap is sent, a syslog message is logged, and the violation counter increments.
- shutdown—A port security violation causes the interface to become error-disabled and to shut down immediately, and the port LED turns off. An SNMP trap is sent, a syslog message is logged, and the violation counter increments. When a secure port is in the error-disabled state, you can bring it out of this state by entering the **errdisable recovery cause** *psecure-violation* global configuration command, or you can manually reenable it by entering the **shutdown** and **no shut down** interface configuration commands. This is the default mode.

- shutdown vlan—Use to set the security violation mode per-VLAN. In this mode, the VLAN is error disabled instead of the entire port when a violation occurs

Table 29-1 Security Violation Mode Actions

Violation Mode	Traffic is Forwarded[1]	Sends SNMP Trap	Sends syslog Message	Displays Error Message[2]	Violation Counter Increments	Shuts Down Port
protect	No	No	No	No	No	No
restrict	No	Yes	Yes	No	Yes	No
shutdown	No	No	No	No	Yes	Yes
shutdown vlan	No	No	Yes	No	Yes	No[3]

1. Packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses.
2. The switch returns an error message if you manually configure an address that would cause a security violation.
3. Shuts down only the VLAN on which the violation occurred.

Default Port Security Configuration

Table 29-2 Default Port Security Configuration

Feature	Default Setting
Port security	Disabled on a port.
Sticky address learning	Disabled.
Maximum number of secure MAC addresses per port	1
Violation mode	Shutdown. The port shuts down when the maximum number of secure MAC addresses is exceeded.
Port security aging	Disabled. Aging time is 0. Static aging is disabled. Type is absolute.

Port Security Configuration Guidelines

- Port security can only be configured on static access ports or trunk ports. A secure port cannot be a dynamic access port.
- A secure port cannot be a destination port for Switched Port Analyzer (SPAN).
- A secure port cannot belong to a Fast EtherChannel port group.

Note Voice VLAN is only supported on access ports and not on trunk ports, even though the configuration is allowed.

- When you enable port security on an interface that is also configured with a voice VLAN, set the maximum allowed secure addresses on the port to two. When the port is connected to a Cisco IP phone, the IP phone requires one MAC address. The Cisco IP phone address is learned on the voice

VLAN, but is not learned on the access VLAN. If you connect a single PC to the Cisco IP phone, no additional MAC addresses are required. If you connect more than one PC to the Cisco IP phone, you must configure enough secure addresses to allow one for each PC and one for the phone.

- When a trunk port configured with port security and assigned to an access VLAN for data traffic and to a voice VLAN for voice traffic, entering the **switchport voice** and **switchport priority extend** interface configuration commands has no effect.

  When a connected device uses the same MAC address to request an IP address for the access VLAN and then an IP address for the voice VLAN, only the access VLAN is assigned an IP address.

- When confguring port security, first specify the total number of MAC addresses you want to allow, by using the **switchport port-security maximum** interface configuration command and then configure the number of access VLANs (**switchport port-security vlan access** interface configuration command) and voice VLANs (**switchport port-security vlan voice** interface configuration command) you want to allow. If you do not specify the total number first, the system returns to the default setting (1 MAC address).
- When you enter a maximum secure address value for an interface, and the new value is greater than the previous value, the new value overwrites the previously configured value. If the new value is less than the previous value and the number of configured secure addresses on the interface exceeds the new value, the command is rejected.
- The switch does not support port security aging of sticky secure MAC addresses.

Table 29-3 Port Security Compatibility with Other Switch Features

Type of Port or Feature on Port	Compatible with Port Security
DTP[1] port[2]	No
Trunk port	Yes
Dynamic-access port[3]	No
Routed port	No
SPAN source port	Yes
SPAN destination port	No
EtherChannel	No
Tunneling port	Yes
Protected port	Yes
IEEE 802.1x port	Yes
Voice VLAN port[4]	Yes
Private VLAN port	Yes
IP source guard	Yes
Dynamic Address Resolution Protocol (ARP) inspection	Yes
FlexLinks	Yes

1. DTP = Dynamic Trunking Protocol
2. A port configured with the **switchport mode dynamic** interface configuration command.
3. A VLAN Query Protocol (VQP) port configured with the **switchport access vlan dynamic** interface configuration command.
4. You must set the maximum allowed secure addresses on the port to two plus the maximum number of secure addresses allowed on the access VLAN.

Port Security Aging

You can use port security aging to set the aging time for all secure addresses on a port. Two types of aging are supported per port:

- Absolute—The secure addresses on the port are deleted after the specified aging time.
- Inactivity—The secure addresses on the port are deleted only if the secure addresses are inactive for the specified aging time.

Use this feature to remove and add devices on a secure port without manually deleting the existing secure MAC addresses and to still limit the number of secure addresses on a port. You can enable or disable the aging of secure addresses on a per-port basis.

Port Security and Private VLANs

Ports that have both port security and private VLANs (PVLANs) configured can be labeled secure PVLAN ports. When a secure address is learned on a secure PVLAN port, the same secure address cannot be learned on another secure PVLAN port belonging to the same primary VLAN. However, an address learned on unsecure PVLAN port can be learned on a secure PVLAN port belonging to same primary VLAN.

Secure addresses that are learned on host port get automatically replicated on associated primary VLANs, and similarly, secure addresses learned on promiscuous ports automatically get replicated on all associated secondary VLANs. Static addresses (using the **mac-address-table static** command) cannot be user configured on a secure port.

Protocol Storm Protection

When a switch is flooded with Address Resolution Protocol (ARP) or control packets, high CPU utilization can cause the CPU to overload. These issues can occur:

- Routing protocol can flap because the protocol control packets are not received, and neighboring adjacencies are dropped.
- Spanning Tree Protocol (STP) reconverges because the STP bridge protocol data unit (BPDU) cannot be sent or received.
- CLI is slow or unresponsive.

Using protocol storm protection, you can control the rate at which control packets are sent to the switch by specifying the upper threshold for the packet flow rate. The supported protocols are ARP, ARP snooping, Dynamic Host Configuration Protocol (DHCP) v4, DHCP snooping, Internet Group Management Protocol (IGMP), and IGMP snooping.

When the packet rate exceeds the defined threshold, the switch drops all traffic arriving on the specified virtual port for 30 seconds. The packet rate is measured again, and protocol storm protection is again applied if necessary.

For further protection, you can manually error disable the virtual port, blocking all incoming traffic on the virtual port. You can manually enable the virtual port or set a time interval for automatic reenabling of the virtual port.

Note Excess packets are dropped on no more than two virtual ports.
Virtual port error disabling is not supported for EtherChannel and Flex Link interfaces.

Protocol storm protection is disabled by default. When it is enabled, auto-recovery of the virtual port is disabled by default.

How to Configure Port-Based Traffic Control

Configuring Storm Control

Configuring Storm Control and Threshold Levels

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface to be configured, and enters interface configuration mode.
Step 3	**storm-control** {**broadcast** | **multicast** | **unicast**} **level** {*level* [*level-low*] | **bps** *bps* [*bps-low*] | **pps** *pps* [*pps-low*]}	Configures broadcast, multicast, or unicast storm control. By default, storm control is disabled. • *level*—Specifies the rising threshold level for broadcast, multicast, or unicast traffic as a percentage (up to two decimal places) of the bandwidth. The port blocks traffic when the rising threshold is reached. The range is 0.00 to 100.00. • (Optional) *level-low*—Specifies the falling threshold level as a percentage (up to two decimal places) of the bandwidth. This value must be less than or equal to the rising suppression value. The port forwards traffic when traffic drops below this level. If you do not configure a falling suppression level, it is set to the rising suppression level. The range is 0.00 to 100.00. If you set the threshold to the maximum value (100 percent), no limit is placed on the traffic. If you set the threshold to 0.0, all broadcast, multicast, and unicast traffic on that port is blocked. • **bps** *bps*—Specifies the rising threshold level for broadcast, multicast, or unicast traffic in bits per second (up to one decimal place). The port blocks traffic when the rising threshold is reached. The range is 0.0 to 10000000000.0. • (Optional) *bps-low*—Specifies the falling threshold level in bits per second (up to one decimal place). It can be less than or equal to the rising threshold level. The port forwards traffic when traffic drops below this level. The range is 0.0 to 10000000000.0.

	Command	Purpose
		• **pps** *pps*—Specifies the rising threshold level for broadcast, multicast, or unicast traffic in packets per second (up to one decimal place). The port blocks traffic when the rising threshold is reached. The range is 0.0 to 10000000000.0. • (Optional) *pps-low*—Specifies the falling threshold level in packets per second (up to one decimal place). It can be less than or equal to the rising threshold level. The port forwards traffic when traffic drops below this level. The range is **0.0 to** 10000000000.0. For BPS and PPS settings, you can use metric suffixes such as k, m, and g for large number thresholds.
Step 4	**storm-control action** {**shutdown** | **trap**}	Specifies the action to be taken when a storm is detected. The default is to filter out the traffic and not to send traps. • **shutdown**—Error-disables the port during a storm. • **trap**—Generates an SNMP trap when a storm is detected.
Step 5	**end**	Returns to privileged EXEC mode.

Configuring Small-Frame Arrival Rate

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**errdisable detect cause small-frame**	Enables the small-frame rate-arrival feature on the switch.
Step 3	**errdisable recovery interval** *interval*	(Optional) Specifies the time to recover from the specified error-disabled state.
Step 4	**errdisable recovery cause small-frame**	(Optional) Configures the recovery time for error-disabled ports to be automatically reenabled after they are error disabled by the arrival of small frames
Step 5	**interface** *interface-id*	Enters interface configuration mode, and specifies the interface to be configured.
Step 6	**small violation-rate** *pps*	Configures the threshold rate for the interface to drop incoming packets and error disable the port. The range is 1 to 10,000 packets per second (pps).
Step 7	**end**	Returns to privileged EXEC mode.

Configuring Protected Ports

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface to be configured, and enter interface configuration mode.

	Command	Purpose
Step 3	**switchport protected**	Configures the interface to be a protected port.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring Port Blocking

Blocking Flooded Traffic on an Interface

Note The interface can be a physical interface or an EtherChannel group. When you block multicast or unicast traffic for a port channel, it is blocked on all ports in the port-channel group.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface to be configured, and enters interface configuration mode.
Step 3	**switchport block multicast**	Blocks unknown multicast forwarding out of the port. **Note** Only pure Layer 2 multicast traffic is blocked. Multicast packets that contain IPv4 or IPv6 information in the header are not blocked.
Step 4	**switchport block unicast**	Blocks unknown unicast forwarding out of the port.
Step 5	**end**	Returns to privileged EXEC mode.

Configuring Port Security

Enabling and Configuring Port Security

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface to be configured, and enters interface configuration mode.
Step 3	**switchport mode** {**access** | **trunk**}	Sets the interface switchport mode as access or trunk. An interface in the default mode (dynamic auto) cannot be configured as a secure port.
Step 4	**switchport voice vlan** *vlan-id*	Enables voice VLAN on a port. *vlan-id*—Specifies the VLAN to be used for voice traffic.
Step 5	**switchport port-security**	Enables port security on the interface.

	Command	Purpose
Step 6	**switchport port-security** [**maximum** *value* [**vlan** {*vlan-list* | {**access** | **voice**}}]]	(Optional) **maximum**—Specifies the maximum number of secure MAC addresses on the port. By default only 1 MAC address is allowed. The maximum number of secure MAC addresses that you can configure on a switch is set by the maximum number of available MAC addresses allowed in the system. This number is set by the active Switch Database Management (SDM) template. See Chapter 11, "Configuring SDM Templates." This number is the total of available MAC addresses, including those used for other Layer 2 functions and any other secure MAC addresses configured on interfaces. (Optional) **vlan**—Sets a per-VLAN maximum value. Enter one of these options after you enter the **vlan** keyword: • *vlan-list*—On a trunk port, sets a per-VLAN maximum value on a range of VLANs separated by a hyphen or a series of VLANs separated by commas. For nonspecified VLANs, the per-VLAN maximum value is used. • **access**—On an access port, specifies the VLAN as an access VLAN. • **voice**—On an access port, specifies the VLAN as a voice VLAN. **Note** The **voice** keyword is available only if a voice VLAN is configured on a port and if that port is not the access VLAN. If an interface is configured for voice VLAN, configure a maximum of two secure MAC addresses.

	Command	Purpose
Step 7	**switchport port-security [violation {protect | restrict | shutdown | shutdown vlan}]**	(Optional) Sets the violation mode, the action to be taken when a security violation is detected, as one of these: • **protect**—When the number of port secure MAC addresses reaches the maximum limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses to drop below the maximum value or increase the number of maximum allowable addresses. You are not notified that a security violation has occurred. **Note** We do not recommend configuring the protect mode on a trunk port. The protect mode disables learning when any VLAN reaches its maximum limit, even if the port has not reached its maximum limit. • **restrict**—When the number of secure MAC addresses reaches the limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses or increase the number of maximum allowable addresses. An SNMP trap is sent, a syslog message is logged, and the violation counter increments. • **shutdown**—The interface is error-disabled when a violation occurs, and the port LED turns off. An SNMP trap is sent, a syslog message is logged, and the violation counter increments. • **shutdown vlan**—Sets the security violation mode per VLAN. In this mode, the VLAN is error disabled instead of the entire port when a violation occurs. **Note** When a secure port is in the error-disabled state, you can bring it out of this state by entering the **errdisable recovery cause psecure-violation** global configuration command. You can manually reenable it by entering the **shutdown** and **no shutdown** interface configuration commands or by using the **clear errdisable interface vlan** privileged EXEC command.

	Command	Purpose
Step 8	**switchport port-security** [**mac-address** *mac-address* [**vlan** {*vlan-id* | {**access** | **voice**}}]	(Optional) Enters a secure MAC address for the interface. You can use this command to enter the maximum number of secure MAC addresses. If you configure fewer secure MAC addresses than the maximum, the remaining MAC addresses are dynamically learned. **Note** If you enable sticky learning after you enter this command, the secure addresses that were dynamically learned are converted to sticky secure MAC addresses and are added to the running configuration. (Optional) **vlan**—Sets a per-VLAN maximum value. Enter one of these options after you enter the **vlan** keyword: • *vlan-id*—On a trunk port, specifies the VLAN ID and the MAC address. If you do not specify a VLAN ID, the native VLAN is used. • **access**—On an access port, specifies the VLAN as an access VLAN. • **voice**—On an access port, specifies the VLAN as a voice VLAN. **Note** The **voice** keyword is available only if a voice VLAN is configured on a port and if that port is not the access VLAN. If an interface is configured for voice VLAN, configure a maximum of two secure MAC addresses.
Step 9	**switchport port-security mac-address sticky**	(Optional) Enables sticky learning on the interface.
Step 10	**switchport port-security mac-address sticky** [*mac-address* | **vlan** {*vlan-id* | {**access** | **voice**}}]	(Optional) Enters a sticky secure MAC address, repeating the command as many times as necessary. If you configure fewer secure MAC addresses than the maximum, the remaining MAC addresses are dynamically learned, are converted to sticky secure MAC addresses, and are added to the running configuration. **Note** If you do not enable sticky learning before this command is entered, an error message appears, and you cannot enter a sticky secure MAC address. (Optional) **vlan**—Sets a per-VLAN maximum value. Enter one of these options after you enter the **vlan** keyword: • *vlan-id*—On a trunk port, specifies the VLAN ID and the MAC address. If you do not specify a VLAN ID, the native VLAN is used. • **access**—On an access port, specifies the VLAN as an access VLAN. • **voice**—On an access port, specifies the VLAN as a voice VLAN. **Note** The **voice** keyword is available only if a voice VLAN is configured on a port and if that port is not the access VLAN.
Step 11	**end**	Returns to privileged EXEC mode.

Enabling and Configuring Port Security Aging

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface to be configured, and enters interface configuration mode.
Step 3	**switchport port-security aging** {**static** | **time** *time* | **type** {**absolute** | **inactivity**}}	Enables or disables static aging for the secure port, or sets the aging time or type. **Note** The switch does not support port security aging of sticky secure addresses. **static**—Enables aging for statically configured secure addresses on this port. **time**—Specifies the aging time for this port. The valid range is from 0 to 1440 minutes. **type**—Specifies the aging type as either absolute or inactivity. • **absolute**—All the secure addresses on this port age out exactly after the time *(*minutes) specified lapses and are removed from the secure address list. • **inactivity**—The secure addresses on this port age out only if there is no data traffic from the secure source addresses for the specified time period.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring Protocol Storm Protection

Enabling Protocol Storm Protection

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**psp** {**arp** | **dhcp** | **igmp**} **pps** *value*	Configures protocol storm protection for ARP, IGMP, or DHCP. *value*—Specifies the threshold value for the number of packets per second. If the traffic exceeds this value, protocol storm protection is enforced. The range is from 5 to 50 packets per second.
Step 3	**errdisable detect cause psp**	(Optional) Enables error-disable detection for protocol storm protection. If this feature is enabled, the virtual port is error-disabled. If this feature is disabled, the port drops excess packets without error-disabling the port.

	Command	Purpose
Step 4	**errdisable recovery interval** *time*	(Optional) Configures an auto-recovery time (in seconds) for error-disabled virtual ports. When a virtual port is error-disabled, the switch auto-recovers after this time. The range is from 30 to 86400 seconds.
Step 5	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining Port-Based Traffic Control

Command	Purpose
show interfaces [*interface-id*] **switchport**	Displays the administrative and operational status of all switching (nonrouting) ports or the specified port, including port blocking and port protection settings.
show storm-control [*interface-id*] [**broadcast** | **multicast** | **unicast**]	Displays storm control suppression levels set on all interfaces or the specified interface for the specified traffic type or for broadcast traffic if no traffic type is entered.
show port-security [**interface** *interface-id*]	Displays port security settings for the switch or for the specified interface, including the maximum allowed number of secure MAC addresses for each interface, the number of secure MAC addresses on the interface, the number of security violations that have occurred, and the violation mode.
show port-security [**interface** *interface-id*] **address**	Displays all secure MAC addresses configured on all switch interfaces or on a specified interface with aging information for each address.
show port-security interface *interface-id* **vlan**	Displays the number of secure MAC addresses configured per VLAN on the specified interface.
show storm-control [*interface-id*] [**broadcast** | **multicast** | **unicast**]	Displays the storm control suppression levels set on the interface for the specified traffic type. If you do not enter a traffic type, broadcast storm control settings are displayed.
show interfaces *interface-id*	Displays the interface configuration.
show interfaces *interface-id* **switchport**	Displays switch-port information.
show port-security	Displays port-security settings for an interface or for the switch.
show psp config {**arp** | **dhcp** | **igmp**}	Displays PSP configuration details for protocols.

Configuration Examples for Port-Based Traffic Control

Enabling Unicast Storm Control: Example

This example shows how to enable unicast storm control on a port with an 87-percent rising suppression level and a 65-percent falling suppression level:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# storm-control unicast level 87 65
```

Enabling Broadcast Address Storm Control on a Port: Example

This example shows how to enable broadcast address storm control on a port to a level of 20 percent. When the broadcast traffic exceeds the configured level of 20 percent of the total available bandwidth of the port within the traffic-storm-control interval, the switch drops all broadcast traffic until the end of the traffic-storm-control interval:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# storm-control broadcast level 20
```

Enabling Small-Frame Arrival Rate: Example

This example shows how to enable the small-frame arrival-rate feature, configure the port recovery time, and configure the threshold for error-disabling a port:

```
Switch# configure terminal
Switch# errdisable detect cause small-frame
Switch# errdisable recovery cause small-frame
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# small-frame violation rate 10000
Switch(config-if)# end
```

Configuring a Protected Port: Example

This example shows how to configure a port as a protected port:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# switchport protected
Switch(config-if)# end
```

Blocking Flooding on a Port: Example

This example shows how to block unicast and Layer 2 multicast flooding on a port:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# switchport block multicast
Switch(config-if)# switchport block unicast
Switch(config-if)# end
```

Configuring Port Security: Examples

This example shows how to enable port security on a port and to set the maximum number of secure addresses to 50. The violation mode is the default, no static secure MAC addresses are configured, and sticky learning is enabled:

```
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# switchport mode access
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security maximum 50
Switch(config-if)# switchport port-security mac-address sticky
```

This example shows how to configure a static secure MAC address on VLAN 3 on a port:

```
Switch(config)# interface gigabitethernet1/2
Switch(config-if)# switchport mode trunk
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security mac-address 0000.02000.0004 vlan 3
```

This example shows how to enable sticky port security on a port, to manually configure MAC addresses for data VLAN and voice VLAN, and to set the total maximum number of secure addresses to 20 (10 for data VLAN and 10 for voice VLAN).

```
Switch(config)# interface FastEthernet1/1
Switch(config-if)# switchport access vlan 21
Switch(config-if)# switchport mode access
Switch(config-if)# switchport voice vlan 22
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security maximum 20
Switch(config-if)# switchport port-security violation restrict
Switch(config-if)# switchport port-security mac-address sticky
Switch(config-if)# switchport port-security mac-address sticky 0000.0000.0002
Switch(config-if)# switchport port-security mac-address 0000.0000.0003
Switch(config-if)# switchport port-security mac-address sticky 0000.0000.0001 vlan voice
Switch(config-if)# switchport port-security mac-address 0000.0000.0004 vlan voice
Switch(config-if)# switchport port-security maximum 10 vlan access
Switch(config-if)# switchport port-security maximum 10 vlan voice
```

Configuring Port Security Aging: Examples

This example shows how to set the aging time as 2 hours for the secure addresses on a port:

```
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# switchport port-security aging time 120
```

This example shows how to set the aging time as 2 minutes for the inactivity aging type with aging enabled for the configured secure addresses on the interface:

```
Switch(config-if)# switchport port-security aging time 2
Switch(config-if)# switchport port-security aging type inactivity
Switch(config-if)# switchport port-security aging static
```

You can verify the previous commands by entering the **show port-security interface** *interface-id* privileged EXEC command.

Configuring Protocol Storm Protection: Example

This example shows how to configure protocol storm protection to drop incoming DHCP traffic on DHCP when it exceeds 35 packets per second:

```
Switch# configure terminal
Switch(config)# psp dhcp pps 35
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 30

Configuring SPAN and RSPAN

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for SPAN and RSPAN

- You must globally configure the **ip device tracking maximum** *limit-number* interface configuration command globally for IPSG for static hosts to work. If you only configure this command on a port without enabling IP device tracking globally or setting an IP device tracking maximum on that interface, IPSG with static hosts will reject all the IP traffic from that interface. This requirement also applies to IPSG with static hosts on a Layer 2 access port.

Restrictions for SPAN and RSPAN

- To use the RSPAN feature, the switch must be running the LAN Base image.
- SPAN for intrusion detection is not supported on the LAN Lite image.
- Two SPAN sessions are supported when the switch is running the LAN Base image.
- One SPAN session is supported when the switch is running the LAN Lite image.

Information About SPAN and RSPAN

SPAN and RSPAN

You can analyze network traffic passing through ports or VLANs by using Switched Port Analyzer (SPAN) or Remote SPAN (RSPAN) to send a copy of the traffic to another port on the switch or on another switch that has been connected to a network analyzer or other monitoring or security device.

SPAN copies (or mirrors) traffic received or sent (or both) on source ports or source VLANs to a destination port for analysis. SPAN does not affect the switching of network traffic on the source ports or VLANs. You must dedicate the destination port for SPAN use. Except for traffic that is required for the SPAN or RSPAN session, destination ports do not receive or forward traffic.

Only traffic that enters or leaves source ports or traffic that enters or leaves source VLANs can be monitored by using SPAN; traffic routed to a source VLAN cannot be monitored. For example, if incoming traffic is being monitored, traffic that gets routed from another VLAN to the source VLAN cannot be monitored; however, traffic that is received on the source VLAN and routed to another VLAN can be monitored.

You can use the SPAN or RSPAN destination port to inject traffic from a network security device. For example, if you connect a Cisco Intrusion Detection System (IDS) sensor appliance to a destination port, the IDS device can send TCP reset packets to close down the TCP session of a suspected attacker.

Local SPAN

Local SPAN supports a SPAN session entirely within one switch; all source ports or source VLANs and destination ports are in the same switch. Local SPAN copies traffic from one or more source ports in any VLAN or from one or more VLANs to a destination port for analysis. For example, in Figure 30-1, all traffic on port 5 (the source port) is mirrored to port 10 (the destination port). A network analyzer on port 10 receives all network traffic from port 5 without being physically attached to port 5.

Figure 30-1 Example of Local SPAN Configuration on a Single Switch

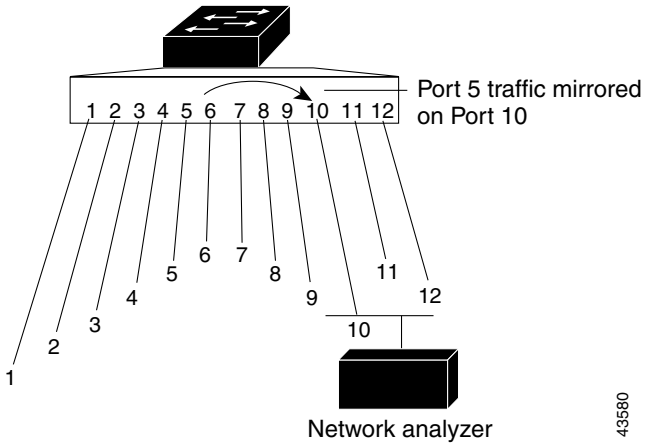

Remote SPAN

RSPAN supports source ports, source VLANs, and destination ports on different switches, enabling remote monitoring of multiple switches across your network. Figure 30-2 shows source ports on Switch A and Switch B. The traffic for each RSPAN session is carried over a user-specified RSPAN VLAN that is dedicated for that RSPAN session in all participating switches. The RSPAN traffic from the source ports or VLANs is copied into the RSPAN VLAN and forwarded over trunk ports carrying the RSPAN VLAN to a destination session monitoring the RSPAN VLAN. Each RSPAN source switch must have either ports or VLANs as RSPAN sources. The destination is always a physical port, as shown on Switch C in the figure.

Figure 30-2 Example of RSPAN Configuration

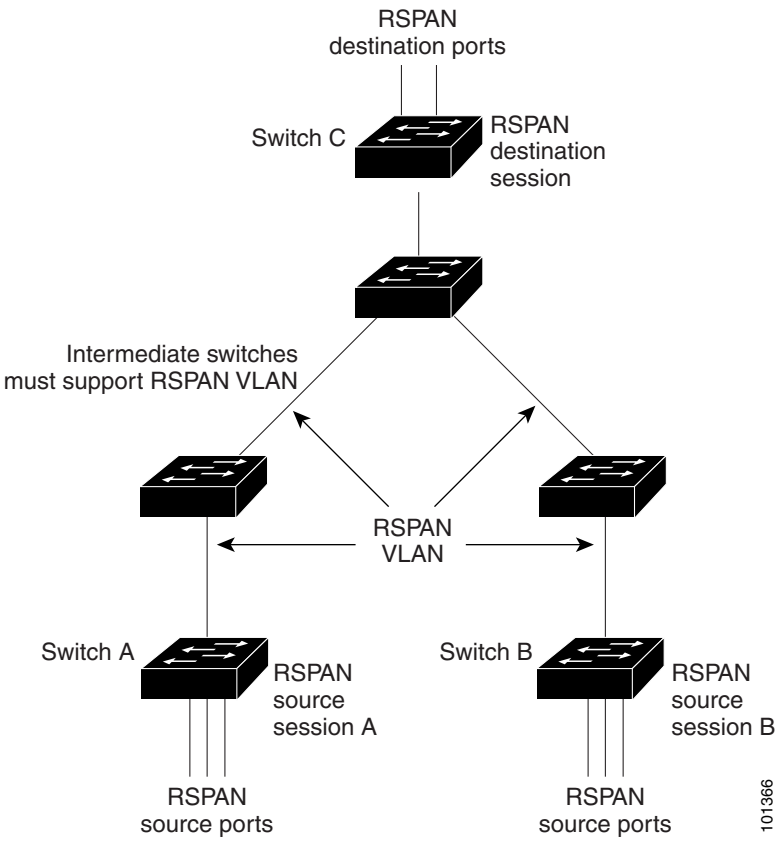

SPAN Sessions

SPAN sessions (local or remote) allow you to monitor traffic on one or more ports, or one or more VLANs, and send the monitored traffic to one or more destination ports.

A local SPAN session is an association of a destination port with source ports or source VLANs, all on a single network device. Local SPAN does not have separate source and destination sessions. Local SPAN sessions gather a set of ingress and egress packets specified by the user and form them into a stream of SPAN data, which is directed to the destination port.

RSPAN consists of at least one RSPAN source session, an RSPAN VLAN, and at least one RSPAN destination session. You separately configure RSPAN source sessions and RSPAN destination sessions on different network devices. To configure an RSPAN source session on a device, you associate a set of source ports or source VLANs with an RSPAN VLAN. The output of this session is the stream of SPAN packets that are sent to the RSPAN VLAN. To configure an RSPAN destination session on another device, you associate the destination port with the RSPAN VLAN. The destination session collects all RSPAN VLAN traffic and sends it out the RSPAN destination port.

An RSPAN source session is very similar to a local SPAN session, except for where the packet stream is directed. In an RSPAN source session, SPAN packets are relabeled with the RSPAN VLAN ID and directed over normal trunk ports to the destination switch.

An RSPAN destination session takes all packets received on the RSPAN VLAN, strips off the VLAN tagging, and presents them on the destination port. Its purpose is to present a copy of all RSPAN VLAN packets (except Layer 2 control packets) to the user for analysis.

There can be more than one source session and more than one destination session active in the same RSPAN VLAN. There can also be intermediate switches separating the RSPAN source and destination sessions. These switches need not be capable of running RSPAN, but they must respond to the requirements of the RSPAN VLAN (see the "RSPAN VLAN" section on page 30-7).

Traffic monitoring in a SPAN session has these restrictions:

- Sources can be ports or VLANs, but you cannot mix source ports and source VLANs in the same session.
- The switch supports up to two source sessions (local SPAN and RSPAN source sessions). You can run both a local SPAN and an RSPAN source session in the same switch. The switch supports a total of 66 source and RSPAN destination sessions.
- You can have multiple destination ports in a SPAN session, but no more than 64 destination ports.
- You can configure two separate SPAN or RSPAN source sessions with separate or overlapping sets of SPAN source ports and VLANs. Both switched and routed ports can be configured as SPAN sources and destinations.
- SPAN sessions do not interfere with the normal operation of the switch. However, an oversubscribed SPAN destination, for example, a 10-Mb/s port monitoring a 100-Mb/s port, can result in dropped or lost packets.
- When RSPAN is enabled, each packet being monitored is transmitted twice, once as normal traffic and once as a monitored packet. Therefore monitoring a large number of ports or VLANs could potentially generate large amounts of network traffic.
- You can configure SPAN sessions on disabled ports; however, a SPAN session does not become active unless you enable the destination port and at least one source port or VLAN for that session.
- The switch does not support a combination of local SPAN and RSPAN in a single session. That is, an RSPAN source session cannot have a local destination port, an RSPAN destination session cannot have a local source port, and an RSPAN destination session and an RSPAN source session that are using the same RSPAN VLAN cannot run on the same switch.

Monitored Traffic Types for SPAN Sessions

- Receive (Rx) SPAN—The goal of receive (or ingress) SPAN is to monitor as much as possible all the packets received by the source interface or VLAN before any modification or processing is performed by the switch. A copy of each packet received by the source is sent to the destination port for that SPAN session.

  Packets that are modified because of routing or quality of service (QoS)—for example, modified Differentiated Services Code Point (DSCP)—are copied before modification.

  Features that can cause a packet to be dropped during receive processing have no effect on ingress SPAN; the destination port receives a copy of the packet even if the actual incoming packet is dropped. These features include IP standard and extended input access control lists (ACLs), ingress QoS policing, VLAN ACLs, and egress QoS policing.
- Transmit (Tx) SPAN—The goal of transmit (or egress) SPAN is to monitor as much as possible all the packets sent by the source interface after all modification and processing is performed by the switch. A copy of each packet sent by the source is sent to the destination port for that SPAN session. The copy is provided after the packet is modified.

  Packets that are modified because of routing—for example, with modified time-to-live (TTL), MAC-address, or QoS values—are duplicated (with the modifications) at the destination port.

Features that can cause a packet to be dropped during transmit processing also affect the duplicated copy for SPAN. These features include IP standard and extended output ACLs and egress QoS policing.

- Both—In a SPAN session, you can also monitor a port or VLAN for both received and sent packets. This is the default.

The default configuration for local SPAN session ports is to send all packets untagged. SPAN also does not normally monitor bridge protocol data unit (BPDU) packets and Layer 2 protocols, such as Cisco Discovery Protocol (CDP), VLAN Trunk Protocol (VTP), Dynamic Trunking Protocol (DTP), Spanning Tree Protocol (STP), and Port Aggregation Protocol (PAgP). However, when you enter the **encapsulation replicate** keywords when configuring a destination port, these changes occur:

- Packets are sent on the destination port with the same encapsulation—untagged or IEEE 802.1Q—that they had on the source port.
- Packets of all types, including BPDU and Layer 2 protocol packets, are monitored.

Therefore, a local SPAN session with encapsulation replicate enabled can have a mixture of untagged and IEEE 802.1Q tagged packets appear on the destination port.

Switch congestion can cause packets to be dropped at ingress source ports, egress source ports, or SPAN destination ports. In general, these characteristics are independent of one another. For example:

- A packet might be forwarded normally but dropped from monitoring due to an oversubscribed SPAN destination port.
- An ingress packet might be dropped from normal forwarding, but still appear on the SPAN destination port.
- An egress packet dropped because of switch congestion is also dropped from egress SPAN.

In some SPAN configurations, multiple copies of the same source packet are sent to the SPAN destination port. For example, a bidirectional (both Rx and Tx) SPAN session is configured for the Rx monitor on port A and Tx monitor on port B. If a packet enters the switch through port A and is switched to port B, both incoming and outgoing packets are sent to the destination port. Both packets are the same (unless a Layer-3 rewrite occurs, in which case the packets are different because of the packet modification).

Source Ports

A source port (also called a *monitored port*) is a switched or routed port that you monitor for network traffic analysis. In a local SPAN session or RSPAN source session, you can monitor source ports or VLANs for traffic in one or both directions. The switch supports any number of source ports (up to the maximum number of available ports on the switch) and any number of source VLANs (up to the maximum number of VLANs supported). However, the switch supports a maximum of two sessions (local or RSPAN) with source ports or VLANs, and you cannot mix ports and VLANs in a single session.

A source port has these characteristics:

- It can be monitored in multiple SPAN sessions.
- Each source port can be configured with a direction (ingress, egress, or both) to monitor.
- It can be any port type (for example, EtherChannel, Fast Ethernet, Gigabit Ethernet, and so forth).
- For EtherChannel sources, you can monitor traffic for the entire EtherChannel or individually on a physical port as it participates in the port channel.
- It can be an access port, trunk port, routed port, or voice VLAN port.
- It cannot be a destination port.

- Source ports can be in the same or different VLANs.
- You can monitor multiple source ports in a single session.

Source VLANs

VLAN-based SPAN (VSPAN) is the monitoring of the network traffic in one or more VLANs. The SPAN or RSPAN source interface in VSPAN is a VLAN ID, and traffic is monitored on all the ports for that VLAN.

VSPAN has these characteristics:

- All active ports in the source VLAN are included as source ports and can be monitored in either or both directions.
- On a given port, only traffic on the monitored VLAN is sent to the destination port.
- If a destination port belongs to a source VLAN, it is excluded from the source list and is not monitored.
- If ports are added to or removed from the source VLANs, the traffic on the source VLAN received by those ports is added to or removed from the sources being monitored.
- You cannot use filter VLANs in the same session with VLAN sources.
- You can monitor only Ethernet VLANs.

VLAN Filtering

When you monitor a trunk port as a source port, by default, all VLANs active on the trunk are monitored. You can limit SPAN traffic monitoring on trunk source ports to specific VLANs by using VLAN filtering.

- VLAN filtering applies only to trunk ports or to voice VLAN ports.
- VLAN filtering applies only to port-based sessions and is not allowed in sessions with VLAN sources.
- When a VLAN filter list is specified, only those VLANs in the list are monitored on trunk ports or on voice VLAN access ports.
- SPAN traffic coming from other port types is not affected by VLAN filtering; that is, all VLANs are allowed on other ports.
- VLAN filtering affects only traffic forwarded to the destination SPAN port and does not affect the switching of normal traffic.

Destination Port

Each local SPAN session or RSPAN destination session must have a destination port (also called a *monitoring port*) that receives a copy of traffic from the source ports or VLANs and sends the SPAN packets to the user, usually a network analyzer.

A destination port has these characteristics:

- For a local SPAN session, the destination port must reside on the same switch as the source port. For an RSPAN session, it is located on the switch containing the RSPAN destination session. There is no destination port on a switch running only an RSPAN source session.
- When a port is configured as a SPAN destination port, the configuration overwrites the original port configuration. When the SPAN destination configuration is removed, the port reverts to its previous configuration. If a configuration change is made to the port while it is acting as a SPAN destination port, the change does not take effect until the SPAN destination configuration had been removed.
- If the port was in an EtherChannel group, it is removed from the group while it is a destination port. If it was a routed port, it is no longer a routed port.
- It can be any Ethernet physical port.
- It cannot be a secure port.
- It cannot be a source port.
- It cannot be an EtherChannel group or a VLAN.
- It can participate in only one SPAN session at a time (a destination port in one SPAN session cannot be a destination port for a second SPAN session).
- When it is active, incoming traffic is disabled. The port does not transmit any traffic except that required for the SPAN session. Incoming traffic is never learned or forwarded on a destination port.
- If ingress traffic forwarding is enabled for a network security device, the destination port forwards traffic at Layer 2.
- It does not participate in any of the Layer 2 protocols (STP, VTP, CDP, DTP, PagP).
- A destination port that belongs to a source VLAN of any SPAN session is excluded from the source list and is not monitored.
- The maximum number of destination ports in a switch is 64.

Local SPAN and RSPAN destination ports behave differently regarding VLAN tagging and encapsulation:

- For local SPAN, if the **encapsulation replicate** keywords are specified for the destination port, these packets appear with the original encapsulation (untaggedor IEEE 802.1Q). If these keywords are not specified, packets appear in the untagged format. Therefore, the output of a local SPAN session with **encapsulation replicate** enabled can contain a mixture of untagged or IEEE 802.1Q-tagged packets.
- For RSPAN, the original VLAN ID is lost because it is overwritten by the RSPAN VLAN identification. Therefore, all packets appear on the destination port as untagged.

RSPAN VLAN

The RSPAN VLAN carries SPAN traffic between RSPAN source and destination sessions. It has these special characteristics:

- All traffic in the RSPAN VLAN is always flooded.
- No MAC address learning occurs on the RSPAN VLAN.
- RSPAN VLAN traffic only flows on trunk ports.
- RSPAN VLANs must be configured in VLAN configuration mode by using the **remote-span** VLAN configuration mode command.
- STP can run on RSPAN VLAN trunks but not on SPAN destination ports.

- An RSPAN VLAN cannot be a private-VLAN primary or secondary VLAN.

For VLANs 1 to 1005 that are visible to VLAN Trunking Protocol (VTP), the VLAN ID and its associated RSPAN characteristic are propagated by VTP. If you assign an RSPAN VLAN ID in the extended VLAN range (1006 to 4096), you must manually configure all intermediate switches.

It is normal to have multiple RSPAN VLANs in a network at the same time with each RSPAN VLAN defining a network-wide RSPAN session. That is, multiple RSPAN source sessions anywhere in the network can contribute packets to the RSPAN session. It is also possible to have multiple RSPAN destination sessions throughout the network, monitoring the same RSPAN VLAN and presenting traffic to the user. The RSPAN VLAN ID separates the sessions.

SPAN and RSPAN Interaction with Other Features

- Routing—SPAN does not monitor routed traffic. VSPAN only monitors traffic that enters or exits the switch, not traffic that is routed between VLANs. For example, if a VLAN is being Rx-monitored and the switch routes traffic from another VLAN to the monitored VLAN, that traffic is not monitored and not received on the SPAN destination port.
- STP—A destination port does not participate in STP while its SPAN or RSPAN session is active. The destination port can participate in STP after the SPAN or RSPAN session is disabled. On a source port, SPAN does not affect the STP status. STP can be active on trunk ports carrying an RSPAN VLAN.
- CDP—A SPAN destination port does not participate in CDP while the SPAN session is active. After the SPAN session is disabled, the port again participates in CDP.
- VTP—You can use VTP to prune an RSPAN VLAN between switches.
- VLAN and trunking—You can modify VLAN membership or trunk settings for source or destination ports at any time. However, changes in VLAN membership or trunk settings for a destination port do not take effect until you remove the SPAN destination configuration. Changes in VLAN membership or trunk settings for a source port immediately take effect, and the respective SPAN sessions automatically adjust accordingly.
- EtherChannel—You can configure an EtherChannel group as a source port but not as a SPAN destination port. When a group is configured as a SPAN source, the entire group is monitored.

  If a physical port is added to a monitored EtherChannel group, the new port is added to the SPAN source port list. If a port is removed from a monitored EtherChannel group, it is automatically removed from the source port list.

  A physical port that belongs to an EtherChannel group can be configured as a SPAN source port and still be a part of the EtherChannel. In this case, data from the physical port is monitored as it participates in the EtherChannel. However, if a physical port that belongs to an EtherChannel group is configured as a SPAN destination, it is removed from the group. After the port is removed from the SPAN session, it rejoins the EtherChannel group. Ports removed from an EtherChannel group remain members of the group, but they are in the *inactive* or *suspended* state.

  If a physical port that belongs to an EtherChannel group is a destination port and the EtherChannel group is a source, the port is removed from the EtherChannel group and from the list of monitored ports.
- Multicast traffic can be monitored. For egress and ingress port monitoring, only a single unedited packet is sent to the SPAN destination port. It does not reflect the number of times the multicast packet is sent.
- A private-VLAN port cannot be a SPAN destination port.

- A secure port cannot be a SPAN destination port.

  For SPAN sessions, do not enable port security on ports with monitored egress when ingress forwarding is enabled on the destination port. For RSPAN source sessions, do not enable port security on any ports with monitored egress.

- An IEEE 802.1x port can be a SPAN source port. You can enable IEEE 802.1x on a port that is a SPAN destination port; however, IEEE 802.1x is disabled until the port is removed as a SPAN destination.

  For SPAN sessions, do not enable IEEE 802.1x on ports with monitored egress when ingress forwarding is enabled on the destination port. For RSPAN source sessions, do not enable IEEE 802.1x on any ports that are egress monitored.

Local SPAN Configuration Guidelines

- For SPAN sources, you can monitor traffic for a single port or VLAN or a series or range of ports or VLANs for each session. You cannot mix source ports and source VLANs within a single SPAN session.
- The destination port cannot be a source port; a source port cannot be a destination port.
- You cannot have two SPAN sessions using the same destination port.
- When you configure a switch port as a SPAN destination port, it is no longer a normal switch port; only monitored traffic passes through the SPAN destination port.
- Entering SPAN configuration commands does not remove previously configured SPAN parameters. You must enter the **no monitor session** {*session_number* | **all** | **local** | **remote**} global configuration command to delete configured SPAN parameters.
- For local SPAN, outgoing packets through the SPAN destination port carry the original encapsulation headers—untagged or IEEE 802.1Q—if the **encapsulation replicate** keywords are specified. If the keywords are not specified, the packets are sent in native form. For RSPAN destination ports, outgoing packets are not tagged.
- You can configure a disabled port to be a source or destination port, but the SPAN function does not start until the destination port and at least one source port or source VLAN are enabled.
- You can limit SPAN traffic to specific VLANs by using the **filter vlan** keyword. If a trunk port is being monitored, only traffic on the VLANs specified with this keyword is monitored. By default, all VLANs are monitored on a trunk port.
- You cannot mix source VLANs and filter VLANs within a single SPAN session.

RSPAN Configuration Guidelines

- All the items in the "Local SPAN Configuration Guidelines" section on page 30-9 apply to RSPAN.
- Because RSPAN VLANs have special properties, you should reserve a few VLANs across your network for use as RSPAN VLANs; do not assign access ports to these VLANs.
- You can apply an output ACL to RSPAN traffic to selectively filter or monitor specific packets. Specify these ACLs on the RSPAN VLAN in the RSPAN source switches.
- For RSPAN configuration, you can distribute the source ports and the destination ports across multiple switches in your network.
- RSPAN does not support BPDU packet monitoring or other Layer 2 switch protocols.

- The RSPAN VLAN is configured only on trunk ports and not on access ports. To avoid unwanted traffic in RSPAN VLANs, make sure that the VLAN remote-span feature is supported in all the participating switches.
- Access ports (including voice VLAN ports) on the RSPAN VLAN are put in the inactive state.
- RSPAN VLANs are included as sources for port-based RSPAN sessions when source trunk ports have active RSPAN VLANs. RSPAN VLANs can also be sources in SPAN sessions. However, since the switch does not monitor spanned traffic, it does not support egress spanning of packets on any RSPAN VLAN identified as the destination of an RSPAN source session on the switch.
- You can configure any VLAN as an RSPAN VLAN as long as these conditions are met:
  - The same RSPAN VLAN is used for an RSPAN session in all the switches.
  - All participating switches support RSPAN.
- We recommend that you configure an RSPAN VLAN before you configure an RSPAN source or a destination session.
- If you enable VTP and VTP pruning, RSPAN traffic is pruned in the trunks to prevent the unwanted flooding of RSPAN traffic across the network for VLAN IDs that are lower than 1005.

Default SPAN and RSPAN Settings

Table 30-1 Default SPAN and RSPAN Settings

Feature	Default Setting
SPAN state (SPAN and RSPAN)	Disabled.
Source port traffic to monitor	Both received and sent traffic (**both**).
Encapsulation type (destination port)	Native form (untagged packets).
Ingress forwarding (destination port)	Disabled,
VLAN filtering	On a trunk interface used as a source port, all VLANs are monitored.
RSPAN VLANs	None configured.

How to Configure SPAN and RSPAN

Creating a Local SPAN Session

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**no monitor session** {*session_number* | **all** | **local** | **remote**}	Removes any existing SPAN configuration for the session. *session_number*—The range is 1 to 66. Specify **all** to remove all SPAN sessions, **local** to remove all local sessions, or **remote** to remove all remote SPAN sessions.

	Command	Purpose
Step 3	**monitor session** *session_number* **source** {**interface** *interface-id* | **vlan** *vlan-id*} [**,** | **-**] [**both** | **rx** | **tx**]	Specifies the SPAN session and the source port (monitored port). *session_number*—The range is 1 to 66. *interface-id*—Specifies the source port or source VLAN to monitor. • source *interface-id*—Specifies the source port to monitor. Valid interfaces include physical interfaces and port-channel logical interfaces (**port-channel** *port-channel-number*). Valid port-channel numbers are 1 to 6. • *vlan-id*—Specifies the source VLAN to monitor. The range is 1 to 4096 (excluding the RSPAN VLAN). **Note** A single session can include multiple sources (ports or VLANs), defined in a series of commands, but you cannot combine source ports and source VLANs in one session. (Optional) [**,** | **-**] Specify a series or range of interfaces. Enter a space before and after the comma; enter a space before and after the hyphen. (Optional) Specify the direction of traffic to monitor. If you do not specify a traffic direction, the SPAN monitors both sent and received traffic. • **both**—Monitors both received and sent traffic. This is the default. • **rx**—Monitors received traffic. • **tx**—Monitors sent traffic. **Note** You can use the **monitor session** *session_number* **source** command multiple times to configure multiple source ports.

	Command	Purpose
Step 4	**monitor session** *session_number* **destination** {**interface** *interface-id* [, | -] [**encapsulation replicate**]}	Specifies the SPAN session and the destination port (monitoring port). *session_number*—Specifies the session number entered in step 3. **Note** For local SPAN, you must use the same session number for the source and destination interfaces. • *interface-id*—Specifies the destination port. The destination interface must be a physical port; it cannot be an EtherChannel, and it cannot be a VLAN. • (Optional) [**,** | **-**]—Specifies a series or range of interfaces. Enter a space before and after the comma; enter a space before and after the hyphen. • (Optional) **encapsulation replicate**—Specifies that the destination interface replicates the source interface encapsulation method. If not selected, the default is to send packets in native form (untagged). **Note** You can use **monitor session** *session_number* **destination** command multiple times to configure multiple destination ports.
Step 5	**end**	Returns to privileged EXEC mode.

Creating a Local SPAN Session and Configuring Incoming Traffic

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**no monitor session** {*session_number* | **all** | **local** | **remote**}	Removes any existing SPAN configuration for the session.
Step 3	**monitor session** *session_number* **source** {**interface** *interface-id* | **vlan** *vlan-id*} [**,** | **-**] [**both** | **rx** | **tx**]	Specifies the SPAN session and the source port (monitored port).

	Command	Purpose
Step 4	**monitor session** *session_number* **destination** {**interface** *interface-id* [, | -] [**encapsulation replicate**] [**ingress** {**dot1q vlan** *vlan-id* | **untagged vlan** *vlan-id* | **vlan** *vlan-id*}]}	Specifies the SPAN session, the destination port, the packet encapsulation, and the ingress VLAN and encapsulation. *session_number*—Specifies the session number entered in Step 3. *interface-id*—Specifies the destination port. The destination interface must be a physical port; it cannot be an EtherChannel, and it cannot be a VLAN. (Optional) [**,** | **-**]—Specifies a series or range of interfaces. Enter a space before and after the comma or hyphen. (Optional) **encapsulation replicate**—Specifies that the destination interface replicates the source interface encapsulation method. If not selected, the default is to send packets in native form (untagged). **ingress**—Enables forwarding of incoming traffic on the destination port and specifies the encapsulation type: • **dot1q vlan** *vlan-id*—Accepts incoming packets with IEEE 802.1Q encapsulation with the specified VLAN as the default VLAN. • **untagged vlan** *vlan-id* or **vlan** *vlan-id*—Accepts incoming packets with untagged encapsulation type with the specified VLAN as the default VLAN.
Step 5	**end**	Returns to privileged EXEC mode.

Specifying VLANs to Filter

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**no monitor session** {*session_number* | **all** | **local** | **remote**}	Removes any existing SPAN configuration for the session. *session_number*—The range is 1 to 66. **all**—Removes all SPAN sessions. **local**—Removes all local sessions. **remote—**Removes all remote SPAN sessions.
Step 3	**monitor session** *session_number* **source interface** *interface-id*	Specifies the characteristics of the source port (monitored port) and SPAN session. *session_number*—The range is 1 to 66. *interface-id*—Specifies the source port to monitor. The interface specified must already be configured as a trunk port.

	Command	Purpose
Step 4	**monitor session** *session_number* **filter vlan** *vlan-id* [**,** | **-**]	Limits the SPAN source traffic to specific VLANs. *session_number*—Enters the session number specified in Step 3. *vlan-id*—The range is 1 to 4096. (Optional) Use a comma (**,**) to specify a series of VLANs, or use a hyphen (**-**) to specify a range of VLANs. Enter a space before and after the comma; enter a space before and after the hyphen.
Step 5	**monitor session** *session_number* **destination** {**interface** *interface-id* [, | -] [**encapsulation replicate**]}	Specifies the SPAN session and the destination port (monitoring port). *session_number*—Specifies the session number entered in Step 3. *interface-id*—Specifies the destination port. The destination interface must be a physical port; it cannot be an EtherChannel, and it cannot be a VLAN. (Optional) [**,** | **-**]—Specifies a series or range of interfaces. Enter a space before and after the comma; enter a space before and after the hyphen. (Optional) **encapsulation replicate**—Specifies that the destination interface replicates the source interface encapsulation method. If not selected, the default is to send packets in native form (untagged).
Step 6	**end**	Returns to privileged EXEC mode.

Configuring a VLAN as an RSPAN VLAN

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**vlan** *vlan-id*	Enters a VLAN ID to create a VLAN, or enters the VLAN ID of an existing VLAN, and enter VLAN configuration mode. The range is 2 to 1001 and 1006 to 4096. The RSPAN VLAN cannot be VLAN 1 (the default VLAN) or VLAN IDs 1002 through 1005 (reserved for Token Ring and FDDI VLANs).
Step 3	**remote-span**	Configures the VLAN as an RSPAN VLAN.
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**copy running-config startup-config**	(Optional) Saves the configuration in the configuration file.

Creating an RSPAN Source Session

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**no monitor session** {*session_number* | **all** | **local** | **remote**}	Removes any existing RSPAN configuration for the session. *session_number*—The range is 1 to 66. **all**—Removes all RSPAN sessions **local—**Removes all local sessions **remote—**Removes all remote SPAN sessions.
Step 3	**monitor session** *session_number* **source** {**interface** *interface-id* | **vlan** *vlan-id*} [**,** | **-**] [**both** | **rx** | **tx**]	Specifies the RSPAN session and the source port (monitored port). *session_number*—The range is 1 to 66. Enter a source port or source VLAN for the RSPAN session: • *interface-id*—Specifies the source port to monitor. Valid interfaces include physical interfaces and port-channel logical interfaces (**port-channel** *port-channel-number*). Valid port-channel numbers are 1 to 48. • *vlan-id*—Specifies the source VLAN to monitor. The range is 1 to 4096 (excluding the RSPAN VLAN). A single session can include multiple sources (ports or VLANs), defined in a series of commands, but you cannot combine source ports and source VLANs in one session. (Optional) [**,** | **-**]—Specifies a series or range of interfaces. Enter a space before and after the comma; enter a space before and after the hyphen. (Optional) Specify the direction of traffic to monitor. If you do not specify a traffic direction, the source interface sends both sent and received traffic. • **both**—Monitors both received and sent traffic. • **rx**—Monitors received traffic. • **tx**—Monitors sent traffic.
Step 4	**monitor session** *session_number* **destination remote vlan** *vlan-id*	Specifies the RSPAN session and the destination RSPAN VLAN. *session_number*—Enters the number defined in Step 3. *vlan-id*—Specifies the source RSPAN VLAN to monitor.
Step 5	**end**	Returns to privileged EXEC mode.
Step 6	**show monitor** [**session** *session_number*] **show running-config**	Verifies the configuration.
Step 7	**copy running-config startup-config**	(Optional) Saves the configuration in the configuration file.

Creating an RSPAN Destination Session

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**vlan** *vlan-id*	Enters the VLAN ID of the RSPAN VLAN created from the source switch, and enters VLAN configuration mode. If both switches are participating in VTP and the RSPAN VLAN ID is from 2 to 1005, Steps 2 through 4 are not required because the RSPAN VLAN ID is propagated through the VTP network.
Step 3	**remote-span**	Identifies the VLAN as the RSPAN VLAN.
Step 4	**exit**	Returns to global configuration mode.
Step 5	**no monitor session** {*session_number* | **all** | **local** | **remote**}	Removes any existing RSPAN configuration for the session. *session_number*—The range is 1 to 66. **all—**Removes all RSPAN sessions **local**—Removes all local sessions **remote—**Removes all remote SPAN sessions.
Step 6	**monitor session** *session_number* **source remote vlan** *vlan-id*	Specifies the RSPAN session and the source RSPAN VLAN. *session_number*—The range is 1 to 66. *vlan-id*—Specifies the source RSPAN VLAN to monitor.
Step 7	**monitor session** *session_number* **destination interface** *interface-id*	Specifies the RSPAN session and the destination interface. *session_number*—Enters the number defined in Step 6. In an RSPAN destination session, you must use the same session number for the source RSPAN VLAN and the destination port. *interface-id*—Specifies the destination interface. The destination interface must be a physical interface. Though visible in the command-line help string, **encapsulation replicate** is not supported for RSPAN. The original VLAN ID is overwritten by the RSPAN VLAN ID, and all packets appear on the destination port as untagged.
Step 8	**end**	Returns to privileged EXEC mode.

Creating an RSPAN Destination Session and Configuring Incoming Traffic

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**no monitor session** {*session_number* | **all** | **local** | **remote**}	Removes any existing SPAN configuration for the session.

	Command	Purpose
Step 3	**monitor session** *session_number* **source remote vlan** *vlan-id*	Specifies the RSPAN session and the source RSPAN VLAN. *session_number*—The range is 1 to 66. *vlan-id*—Specifies the source RSPAN VLAN to monitor.
Step 4	**monitor session** *session_number* **destination** {**interface** *interface-id* [, | -] [**ingress** {**dot1q vlan** *vlan-id* | **untagged vlan** *vlan-id* | **vlan** *vlan-id*}]}	Specifies the SPAN session, the destination port, the packet encapsulation, and the incoming VLAN and encapsulation. *session_number*—Enters the number defined in Step 4. In an RSPAN destination session, you must use the same session number for the source RSPAN VLAN and the destination port. *interface-id*—Specifies the destination interface. The destination interface must be a physical interface. Though visible in the command-line help string, **encapsulation replicate** is not supported for RSPAN. The original VLAN ID is overwritten by the RSPAN VLAN ID, and all packets appear on the destination port as untagged. (Optional) [**,** | **-**]—Specifies a series or range of interfaces. Enter a space before and after the comma; enter a space before and after the hyphen. Enter **ingress** with additional keywords to enable forwarding of incoming traffic on the destination port and to specify the encapsulation type: • **dot1q vlan** *vlan-id*—Forwards incoming packets with IEEE 802.1Q encapsulation with the specified VLAN as the default VLAN. • **untagged vlan** *vlan-id* or **vlan** *vlan-id*—Forwards incoming packets with untagged encapsulation type with the specified VLAN as the default VLAN.
Step 5	**end**	Returns to privileged EXEC mode.

Specifying VLANs to Filter

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**no monitor session** {*session_number* | **all** | **local** | **remote**}	Removes any existing SPAN configuration for the session. *session_number*—The range is 1 to 66. **all—**Removes all SPAN sessions. **local—**Removes all local sessions. **remote**—Removes all remote SPAN sessions.

	Command	Purpose
Step 3	**monitor session** *session_number* **source interface** *interface-id*	Specifies the characteristics of the source port (monitored port) and SPAN session. *session_number*—The range is 1 to 66. *interface-id*—Specifies the source port to monitor. The interface specified must already be configured as a trunk port.
Step 4	**monitor session** *session_number* **filter vlan** *vlan-id* [**,** | **-**]	Limits the SPAN source traffic to specific VLANs. *session_number*—Enters the session number specified in step 3. *vlan-id*—The range is 1 to 4096. (Optional) Use a comma (**,**) to specify a series of VLANs or use a hyphen (**-**) to specify a range of VLANs. Enter a space before and after the comma; enter a space before and after the hyphen.
Step 5	**monitor session** *session_number* **destination remote vlan** *vlan-id*	Specifies the RSPAN session and the destination remote VLAN (RSPAN VLAN). *session_number*—Enter the session number specified in step 3. *vlan-id*—Specifies the RSPAN VLAN to carry the monitored traffic to the destination port.
Step 6	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining SPAN and RSPAN

show monitor [**session** *session_number*]	Verifies the SPAN or RSPAN configuration.

Configuration Examples for SPAN and RSPAN

Configuring a Local SPAN Session: Example

This example shows how to set up SPAN session 1 for monitoring source port traffic to a destination port. First, any existing SPAN configuration for session 1 is deleted, and then bidirectional traffic is mirrored from source Gigabit Ethernet port 1 to destination Gigabit Ethernet port 2, retaining the encapsulation method.

```
Switch(config)# no monitor session 1
Switch(config)# monitor session 1 source interface gigabitethernet1/1
Switch(config)# monitor session 1 destination interface gigabitethernet1/2
encapsulation replicate
Switch(config)# end
```

Modifying Local SPAN Sessions: Examples

This example shows how to remove port 1 as a SPAN source for SPAN session 1:

```
Switch(config)# no monitor session 1 source interface gigabitethernet1/1
```

```
Switch(config)# end
```

This example shows how to disable received traffic monitoring on port 1, which was configured for bidirectional monitoring:

```
Switch(config)# no monitor session 1 source interface gigabitethernet1/1 rx
```

The monitoring of traffic received on port 1 is disabled, but traffic sent from this port continues to be monitored.

This example shows how to remove any existing configuration on SPAN session 2, configure SPAN session 2 to monitor received traffic on all ports belonging to VLANs 1 through 3, and send it to destination Gigabit Ethernet port 2. The configuration is then modified to also monitor all traffic on all ports belonging to VLAN 10.

```
Switch(config)# no monitor session 2
Switch(config)# monitor session 2 source vlan 1 - 3 rx
Switch(config)# monitor session 2 destination interface gigabitethernet1/2
Switch(config)# monitor session 2 source vlan 10
Switch(config)# end
```

This example shows how to remove any existing configuration on SPAN session 2, configure SPAN session 2 to monitor received traffic on Gigabit Ethernet source port 1, and send it to destination Gigabit Ethernet port 2 with the same egress encapsulation type as the source port, and to enable ingress forwarding with IEEE 802.1Q encapsulation and VLAN 6 as the default ingress VLAN.

```
Switch(config)# no monitor session 2
Switch(config)# monitor session 2 source gigabitethernet1/1 rx
Switch(config)# monitor session 2 destination interface gigabitethernet1/2 encapsulation
replicate ingress dot1q vlan 6
Switch(config)# end
```

To monitor all VLANs on the trunk port, use the **no monitor session** *session_number* **filter** global configuration command.

This example shows how to remove any existing configuration on SPAN session 2, configure SPAN session 2 to monitor traffic received on Gigabit Ethernet trunk port 2, and send traffic for only VLANs 1 through 5 and VLAN 9 to destination Gigabit Ethernet port 1:

```
Switch(config)# no monitor session 2
Switch(config)# monitor session 2 source interface gigabitethernet1/2 rx
Switch(config)# monitor session 2 filter vlan 1 - 5, 9
Switch(config)# monitor session 2 destination interface gigabitethernet1/1
Switch(config)# end
```

Configuring an RSPAN: Example

This example shows how to create RSPAN VLAN 901:

```
Switch(config)# vlan 901
Switch(config-vlan)# remote span
Switch(config-vlan)# end
```

Configuring a VLAN for a SPAN Session: Example

This example shows how to configure VLAN 901 as the source remote VLAN and port 1 as the destination interface:

```
Switch(config)# monitor session 1 source remote vlan 901
Switch(config)# monitor session 1 destination interface gigabitethernet1/1
Switch(config)# end
```

Modifying RSPAN Sessions: Examples

This example shows how to remove any existing RSPAN configuration for session 1, configure RSPAN session 1 to monitor multiple source interfaces, and configure the destination as RSPAN VLAN 901:

```
Switch(config)# no monitor session 1
Switch(config)# monitor session 1 source interface gigabitethernet1/1 tx
Switch(config)# monitor session 1 source interface gigabitethernet1/2 rx

Switch(config)# monitor session 1 source interface port-channel 2
Switch(config)# monitor session 1 destination remote vlan 901
Switch(config)# end
```

This example shows how to configure VLAN 901 as the source remote VLAN in RSPAN session 2, to configure Gigabit Ethernet source port 2 as the destination interface, and to enable forwarding of incoming traffic on the interface with VLAN 6 as the default receiving VLAN:

```
Switch(config)# monitor session 2 source remote vlan 901
Switch(config)# monitor session 2 destination interface gigabitethernet1/2 ingress vlan 6
Switch(config)# end
```

This example shows how to remove any existing configuration on RSPAN session 2, configure RSPAN session 2 to monitor traffic received on trunk port 2, and send traffic for only VLANs 1 through 5 and 9 to destination RSPAN VLAN 902:

```
Switch(config)# no monitor session 2
(config)# monitor session 2 source interface gigabitethernet1/2 rx
Switch(config)# monitor session 2 filter vlan 1 - 5, 9
Switch(config)# monitor session 2 destination remote vlan 902
Switch(config)# end
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

CHAPTER 31

Configuring LLDP, LLDP-MED, and Wired Location Service

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Restrictions for LLDP, LLDP-MED, and Wired Location Service

- To use the the following features, the switch must be running the LAN Base image:
  - LLDP-MED location 802.lab
  - LLDP-MED integration for CoS/DSCP
  - Network policy TLV and location TLV
  - Wired location service

Information About LLDP, LLDP-MED, and Wired Location Service

The Cisco Discovery Protocol (CDP) is a device discovery protocol that runs over Layer 2 (the data link layer) on all Cisco-manufactured devices (routers, bridges, access servers, and switches). CDP allows network management applications to automatically discover and learn about other Cisco devices connected to the network.

To support non-Cisco devices and to allow for interoperability between other devices, the switch supports the IEEE 802.1AB Link Layer Discovery Protocol (LLDP). LLDP is a neighbor discovery protocol that is used for network devices to advertise information about themselves to other devices on the network. This protocol runs over the data-link layer, which allows two systems running different network layer protocols to learn about each other.

LLDP supports a set of attributes that it uses to discover neighbor devices. These attributes contain type, length, and value descriptions and are referred to as TLVs. LLDP supported devices can use TLVs to receive and send information to their neighbors. This protocol can advertise details such as configuration information, device capabilities, and device identity.

The switch supports these basic management TLVs. These are mandatory LLDP TLVs.

- Port description TLV
- System name TLV
- System description TLV
- System capabilities TLV
- Management address TLV

These organizationally specific LLDP TLVs are also advertised to support LLDP-MED:

- Port VLAN ID TLV ((IEEE 802.1 organizationally specific TLVs)
- MAC/PHY configuration/status TLV(IEEE 802.3 organizationally specific TLVs)

Note A switch stack appears as a single switch in the network. Therefore, LLDP discovers the switch stack, not the individual stack members.

LLDP-MED

LLDP for Media Endpoint Devices (LLDP-MED) is an extension to LLDP that operates between endpoint devices such as IP phones and network devices such as switches. It specifically provides support for voice over IP (VoIP) applications and provides additional TLVs for capabilities discovery, network policy, Power over Ethernet, inventory management and location information. By default, all LLDP-MED TLVs are enabled.

LLDP-MED supports these TLVs:

- LLDP-MED capabilities TLV

  Allows LLDP-MED endpoints to determine the capabilities that the connected device supports and has enabled.

- Network policy TLV

  Allows both network connectivity devices and endpoints to advertise VLAN configurations and associated Layer 2 and Layer 3 attributes for the specific application on that port. For example, the switch can notify a phone of the VLAN number that it should use. The phone can connect to any switch, obtain its VLAN number, and then start communicating with the call control.

  By defining a network-policy profile TLV, you can create a profile for voice and voice-signalling by specifying the values for VLAN, class of service (CoS), differentiated services code point (DSCP), and tagging mode. These profile attributes are then maintained centrally on the switch and propagated to the phone.

- Power management TLV

  Enables advanced power management between LLDP-MED endpoint and network connectivity devices. Allows switches and phones to convey power information, such as how the device is powered, power priority, and how much power the device needs.

- Inventory management TLV

  Allows an endpoint to send detailed inventory information about itself to the switch, including information hardware revision, firmware version, software version, serial number, manufacturer name, model name, and asset ID TLV.

- Location TLV

  Provides location information from the switch to the endpoint device. The location TLV can send this information:

  – Civic location information

    Provides the civic address information and postal information. Examples of civic location information are street address, road name, and postal community name information.

  – ELIN location information

    Provides the location information of a caller. The location is determined by the emergency location identifier number (ELIN), which is a phone number that routes an emergency call to the local public safety answering point (PSAP) and which the PSAP can use to call back the emergency caller.

Wired Location Service

The switch uses the wired location service feature to send location and attachment tracking information for its connected devices to a Cisco Mobility Services Engine (MSE). The tracked device can be a wireless endpoint, a wired endpoint, or a wired switch or controller. The switch notifies the MSE of device link up and link down events through the Network Mobility Services Protocol (NMSP) location and attachment notifications.

The MSE starts the NMSP connection to the switch, which opens a server port. When the MSE connects to the switch there are a set of message exchanges to establish version compatibility and service exchange information followed by location information synchronization. After connection, the switch periodically sends location and attachment notifications to the MSE. Any link up or link down events detected during an interval are aggregated and sent at the end of the interval.

When the switch determines the presence or absence of a device on a link-up or link-down event, it obtains the client-specific information such as the MAC address, IP address, and username. If the client is LLDP-MED- or CDP-capable, the switch obtains the serial number and UDI through the LLDP-MED location TLV or CDP.

Depending on the device capabilities, the switch obtains this client information at link up:

- Slot and port specified in port connection
- MAC address specified in the client MAC address
- IP address specified in port connection
- 802.1X username if applicable
- Device category is specified as a *wired station*
- State is specified as *new*
- Serial number, UDI
- Model number
- Time in seconds since the switch detected the association

Depending on the device capabilities, the switch obtains this client information at link down:

- Slot and port that was disconnected
- MAC address
- IP address
- 802.1X username if applicable
- Device category is specified as a *wired station*
- State is specified as *delete*
- Serial number, UDI
- Time in seconds since the switch detected the disassociation

When the switch shuts down, it sends an attachment notification with the state *delete* and the IP address before closing the NMSP connection to the MSE. The MSE interprets this notification as disassociation for all the wired clients associated with the switch.

If you change a location address on the switch, the switch sends an NMSP location notification message that identifies the affected ports and the changed address information.

Default LLDP Configuration

Table 31-1 Default LLDP Configuration

Feature	Default Setting
LLDP global state	Disabled.
LLDP holdtime (before discarding)	120 seconds.
LLDP timer (packet update frequency)	30 seconds.
LLDP reinitialization delay	2 seconds.
LLDP tlv-select	Disabled to send and receive all TLVs.
LLDP interface state	Disabled.
LLDP receive	Disabled.
LLDP transmit	Disabled.
LLDP med-tlv-select	Disabled to send all LLDP-MED TLVs. When LLDP is globally enabled, LLDP-MED-TLV is also enabled.

LLDP, LLDP-MED, and Wired Location Service Configuration Guidelines

- If the interface is configured as a tunnel port, LLDP is automatically disabled.
- If you first configure a network-policy profile on an interface, you cannot apply the **switchport voice vlan** command on the interface. If the **switchport voice vlan** *vlan-id* is already configured on an interface, you can apply a network-policy profile on the interface. This way the interface has the voice or voice-signaling VLAN network-policy profile applied on the interface.
- You cannot configure static secure MAC addresses on an interface that has a network-policy profile.
- You cannot configure a network-policy profile on a private-VLAN port.

- For wired location to function, you must first enter the **ip device tracking** global configuration command.

LLDP-MED TLVs

By default, the switch only sends LLDP packets until it receives LLDP-MED packets from the end device. It then sends LLDP packets with MED TLVs. When the LLDP-MED entry has been aged out, it only sends LLDP packets.

By using the **lldp** interface configuration command, you can configure the interface not to send the TLVs listed in this table.

Table 31-2 LLDP-MED TLVs

LLDP-MED TLV	Description
inventory-management	LLDP-MED inventory management TLV
location	LLDP-MED location TLV (only on LAN Base image)
network-policy	LLDP-MED network policy TLV (only on LAN Base image)
power-management	LLDP-MED power management TLV

How to Configure LLDP, LLDP-MED, and Wired Location Service

Enabling LLDP

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**lldp run**	Enables LLDP globally on the switch.
Step 3	**interface** *interface-id*	Specifies the interface on which you are enabling LLDP, and enter interface configuration mode.
Step 4	**lldp transmit**	Enables the interface to send LLDP packets.
Step 5	**lldp receive**	Enables the interface to receive LLDP packets.
Step 6	**end**	Returns to privileged EXEC mode.

Configuring LLDP Characteristics

You can configure the frequency of LLDP updates, the amount of time to hold the information before discarding it, and the initialization delay time. You can also select the LLDP and LLDP-MED TLVs to send and receive.

Note Steps 2 through 5 are optional and can be performed in any order.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**lldp holdtime** *seconds*	(Optional) Specifies the amount of time a receiving device should hold the information from your device before discarding it. The range is 0 to 65535 seconds; the default is 120 seconds.
Step 3	**lldp reinit** *delay*	(Optional) Specifies the delay time in seconds for LLDP to initialize on an interface. The range is 2 to 5 seconds; the default is 2 seconds.
Step 4	**lldp timer** *rate*	(Optional) Sets the sending frequency of LLDP updates in seconds. The range is 5 to 65534 seconds; the default is 30 seconds.
Step 5	**lldp tlv-select**	(Optional) Specifies the LLDP TLVs to send or receive.
Step 6	**lldp med-tlv-select**	(Optional) Specifies the LLDP-MED TLVs to send or receive.
Step 7	**end**	Returns to privileged EXEC mode.

Configuring LLDP-MED TLVs

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface on which you are configuring an LLDP-MED TLV, and enters interface configuration mode.
Step 3	**lldp med-tlv-select** *tlv*	Specifies the TLV to enable.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring Network-Policy TLV

This task explains how to create a network-policy profile, configure the policy attributes, and apply it to an interface.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**network-policy profile** *profile number*	Specifies the network-policy profile number, and enters network-policy configuration mode. The range is 1 to 4294967295.

	Command	Purpose
Step 3	{**voice** | **voice-signaling**} **vlan** [*vlan-id* {**cos** *cvalue* | **dscp** *dvalue*}] | [[**dot1p** {**cos** *cvalue* | **dscp** *dvalue*}] | **none** | **untagged**]	Configures the policy attributes: **voice**—Specifies the voice application type. **voice-signaling**—Specifies the voice-signaling application type. **vlan**—Specifies the native VLAN for voice traffic. *vlan-id*—(Optional) Specifies the VLAN for voice traffic. The range is 1 to 4096. **cos** *cvalue*—(Optional) Specifies the Layer 2 priority class of service (CoS) for the configured VLAN. The range is 0 to 7; the default is 0. **dscp** *dvalue*—(Optional) Specifies the differentiated services code point (DSCP) value for the configured VLAN. The range is 0 to 63; the default is 0. **dot1p**—(Optional) Configures the telephone to use IEEE 802.1p priority tagging and use VLAN 0 (the native VLAN). **none**—(Optional) Does not instruct the IP telephone about the voice VLAN. The telephone uses the configuration from the telephone key pad. **untagged**—(Optional) Configures the telephone to send untagged voice traffic. This is the default for the telephone.
Step 4	**exit**	Returns to global configuration mode.
Step 5	**interface** *interface-id*	Specifies the interface on which you are configuring a network-policy profile, and enter interface configuration mode.
Step 6	**network-policy** *profile number*	Specifies the network-policy profile number.
Step 7	**lldp med-tlv-select network-policy**	Specifies the network-policy TLV.
Step 8	**end**	Returns to privileged EXEC mode.

Configuring Location TLV and Wired Location Service

This task explains how to configure location information for an endpoint and to apply it to an interface.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**location** {**admin-tag** *string* | **civic-location identifier** *id* | **elin-location** *string* **identifier** *id*}	Specifies the location information for an endpoint. • **admin-tag**—Specifies an administrative tag or site information. • **civic-location**—Specifies civic location information. • **elin-location**—Specifies emergency location information (ELIN). • **identifier** *id*—Specifies the ID for the civic location. • *string*—Specifies the site or location information in alphanumeric format.
Step 3	**exit**	Returns to global configuration mode.

	Command	Purpose
Step 4	**interface** *interface-id*	Specifies the interface on which you are configuring the location information, and enters interface configuration mode.
Step 5	**location** {**additional-location-information** *word* | **civic-location-id** *id* | **elin-location-id** *id*}	Enters location information for an interface: **additional-location-information**—Specifies additional information for a location or place. **civic-location-id**—Specifies global civic location information for an interface. **elin-location-id**—Specifies emergency location information for an interface. *id*—Specifies the ID for the civic location or the ELIN location. The ID range is 1 to 4095. *word*—Specifies a word or phrase with additional location information.
Step 6	**end**	Returns to privileged EXEC mode.
Step 7	**nmsp enable**	Enables the NMSP features on the switch.
Step 8	**nmsp notification interval** {**attachment** | **location**} *interval-seconds*	Specifies the NMSP notification interval. **attachment**—Specifies the attachment notification interval. **location—**Specifies the location notification interval. *interval-seconds*—Duration in seconds before the switch sends the MSE the location or attachment updates. The range is 1 to 30; the default is 30.
Step 9	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining LLDP, LLDP-MED, and Wired Location Service

Command	Description
clear lldp counters	Resets the traffic counters to zero.
clear lldp table	Deletes the LLDP neighbor information table.
clear nmsp statistics	Clears the NMSP statistic counters.
show lldp	Displays global information, such as frequency of transmissions, the holdtime for packets being sent, and the delay time before LLDP initializes on an interface.
show lldp entry *entry-name*	Displays information about a specific neighbor. You can enter an asterisk (*) to display all neighbors, or you can enter the neighbor name.
show lldp interface [*interface-id*]	Displays information about interfaces with LLDP enabled. You can limit the display to a specific interface.

Command	Description
show lldp neighbors [*interface-id*] [**detail**]	Displays information about neighbors, including device type, interface type and number, holdtime settings, capabilities, and port ID. You can limit the display to neighbors of a specific interface or expand the display for more detailed information.
show lldp traffic	Displays LLDP counters, including the number of packets sent and received, number of packets discarded, and number of unrecognized TLVs.
show location admin-tag *string*	Displays the location information for the specified administrative tag or site.
show location civic-location identifier *id*	Displays the location information for a specific global civic location.
show location elin-location identifier *id*	Displays the location information for an emergency location.
show network-policy profile	Displays the configured network-policy profiles.
show nmsp	Displays the NMSP information.

Configuration Examples for Configuring LLDP, LLDP-MED, and Wired Location Service

Enabling LLDP: Examples

This example shows how to globally enable LLDP:

```
Switch# configure terminal
Switch(config)# lldp run
Switch(config)# end
```

This example shows how to enable LLDP on an interface:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# lldp transmit
Switch(config-if)# lldp receive
Switch(config-if)# end
```

Configuring LDP Parameters: Examples

This example shows how to configure LLDP parameters:

```
Switch# configure terminal
Switch(config)# lldp holdtime 120
Switch(config)# lldp reinit 2
Switch(config)# lldp timer 30
Switch(config)# end
```

Configuring TLV: Example

This example shows how to enable a TLV on an interface:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/1
```

```
Switch(config-if)# lldp med-tlv-select inventory-management
Switch(config-if)# end
```

Configuring Network Policy: Example

This example shows how to configure VLAN 100 for voice application with CoS and to enable the network-policy profile and network-policy TLV on an interface:

```
Switch# configure terminal
Switch(config)# network-policy profile 1
Switch(config-network-policy)# voice vlan 100 cos 4
Switch(config-network-policy)# exit
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# network-policy profile 1
Switch(config-if)# lldp med-tlv-select network-policy
```

Configuring Voice Application: Example

This example shows how to configure the voice application type for the native VLAN with priority tagging:

```
Switch(config-network-policy)# voice vlan dot1p cos 4
Switch(config-network-policy)# voice vlan dot1p dscp 34
```

Configuring Civic Location Information: Example

This example shows how to configure civic location information on the switch:

```
Switch(config)# location civic-location identifier 1
Switch(config-civic)# number 3550
Switch(config-civic)# primary-road-name "Cisco Way"
Switch(config-civic)# city "San Jose"
Switch(config-civic)# state CA
Switch(config-civic)# building 19
Switch(config-civic)# room C6
Switch(config-civic)# county "Santa Clara"
Switch(config-civic)# country US
Switch(config-civic)# end
```

Enabling NMSP: Example

This example shows how to enable NMSP on a switch and to set the location notification time to 10 seconds:

```
Switch(config)# nmsp enable
Switch(config)# nmsp notification interval location 10
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands Cisco IOS system management commands	*Cisco IOS Configuration Fundamentals Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 32

Configuring CDP

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Information About CDP

CDP

CDP is a device discovery protocol that runs over Layer 2 (the data link layer) on all Cisco-manufactured devices (routers, bridges, access servers, and switches) and allows network management applications to discover Cisco devices that are neighbors of already known devices. With CDP, network management applications can learn the device type and the Simple Network Management Protocol (SNMP) agent address of neighboring devices running lower-layer, transparent protocols. This feature enables applications to send SNMP queries to neighboring devices.

CDP runs on all media that support Subnetwork Access Protocol (SNAP). Because CDP runs over the data-link layer only, two systems that support different network-layer protocols can learn about each other.

Each CDP-configured device sends periodic messages to a multicast address, advertising at least one address at which it can receive SNMP messages. The advertisements also contain time-to-live, or holdtime information, which is the length of time a receiving device holds CDP information before discarding it. Each device also listens to the messages sent by other devices to learn about neighboring devices.

On the switch, CDP enables Network Assistant to display a graphical view of the network. The switch uses CDP to find cluster candidates and maintain information about cluster members and other devices up to three cluster-enabled devices away from the command switch by default.

For a switch and connected endpoint devices running Cisco Medianet, these events occur:

- CDP identifies connected endpoints that communicate directly with the switch.
- Only one wired switch reports the location information to prevent duplicate reports of neighboring devices.
- The wired switch and the endpoints both send and receive location information.

The switch supports CDP Version 2.

Default CDP Configuration

Table 32-1 Default CDP Configuration

Feature	Default Setting
CDP global state	Enabled
CDP interface state	Enabled
CDP timer (packet update frequency)	60 seconds
CDP holdtime (before discarding)	180 seconds
CDP Version-2 advertisements	Enabled

How to Configure CDP

Configuring the CDP Parameters

You can configure the frequency of CDP updates, the amount of time to hold the information before discarding it, and whether or not to send Version-2 advertisements.

Note Steps 2 through 4 are all optional and can be performed in any order.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**cdp timer** *seconds*	(Optional) Sets the transmission frequency of CDP updates in seconds. The range is 5 to 254; the default is 60 seconds.
Step 3	**cdp holdtime** *seconds*	(Optional) Specifies the amount of time a receiving device should hold the information sent by your device before discarding it. The range is 10 to 255 seconds; the default is 180 seconds.
Step 4	**cdp advertise-v2**	(Optional) Configures CDP to send Version-2 advertisements. This is the default state.
Step 5	**end**	Returns to privileged EXEC mode.

Disabling CDP

CDP is enabled by default.

Note Switch clusters and other Cisco devices (such as Cisco IP Phones) regularly exchange CDP messages. Disabling CDP can interrupt cluster discovery and device connectivity.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**no cdp run**	Disables CDP globally.
Step 3	**interface** *interface-id*	Specifies the interface on which you are disabling CDP, and enters interface configuration mode.
Step 4	**no cdp enable**	Disables CDP on the interface.
Step 5	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining CDP

Command	Description
clear cdp counters	Resets the traffic counters to zero.
clear cdp table	Deletes the CDP table of information about neighbors.
show cdp	Displays global information, such as frequency of transmissions and the holdtime for packets being sent.
show cdp entry *entry-name* [**protocol** | **version**]	Displays information about a specific neighbor. You can enter an asterisk (*) to display all CDP neighbors, or you can enter the name of the neighbor about which you want information. You can also limit the display to information about the protocols enabled on the specified neighbor or information about the version of software running on the device.
show cdp interface [*interface-id*]	Displays information about interfaces where CDP is enabled. You can limit the display to the interface about which you want information.
show cdp neighbors [*interface-id*] [**detail**]	Displays information about neighbors, including device type, interface type and number, holdtime settings, capabilities, platform, and port ID. You can limit the display to neighbors of a specific interface or expand the display to provide more detailed information.
show cdp traffic	Displays CDP counters, including the number of packets sent and received and checksum errors.

Configuration Examples for CDP

Configuring CDP Parameters: Example

This example shows how to configure CDP parameters:

```
Switch# configure terminal
Switch(config)# cdp timer 50
Switch(config)# cdp holdtime 120
Switch(config)# cdp advertise-v2
Switch(config)# end
```

Enabling CDP: Examples

This example shows how to enable CDP on a port when it has been disabled:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# cdp enable
Switch(config-if)# end
```

Note Voice VLAN is not counted against port security when CDP is disabled on the switch interface.

This example shows how to enable CDP if it has been disabled:

```
Switch# configure terminal
Switch(config)# cdp run
Switch(config)# end
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands Cisco IOS system management commands	*Cisco IOS Configuration Fundamentals Command Reference*
Switch cluster configuration	Chapter 6, “Configuring Switch Clusters”

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

CHAPTER 33

Configuring UDLD

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for UDLD

- When configuring the mode (normal or aggressive), make sure that the same mode is configured on both sides of the link.

Restrictions for UDLD

- UDLD is not supported on ATM ports.
- A UDLD-capable port cannot detect a unidirectional link if it is connected to a UDLD-incapable port of another switch.
- Loop guard works only on point-to-point links. We recommend that each end of the link has a directly connected device that is running STP.

Information About UDLD

UDLD

UniDirectional Link Detection (UDLD) is a Layer 2 protocol that enables devices connected through fiber-optic or twisted-pair Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists. All connected devices must support UDLD for the protocol to successfully identify and disable unidirectional links. When UDLD detects a unidirectional link, it disables the affected port and alerts you. Unidirectional links can cause a variety of problems, including spanning-tree topology loops.

Modes of Operation

UDLD supports two modes of operation: normal (the default) and aggressive. In normal mode, UDLD can detect unidirectional links due to misconnected ports on fiber-optic connections. In aggressive mode, UDLD can also detect unidirectional links due to one-way traffic on fiber-optic and twisted-pair links and to misconnected ports on fiber-optic links.

In normal and aggressive modes, UDLD works with the Layer 1 mechanisms to learn the physical status of a link. At Layer 1, autonegotiation takes care of physical signaling and fault detection. UDLD performs tasks that autonegotiation cannot perform, such as detecting the identities of neighbors and shutting down misconnected ports. When you enable both autonegotiation and UDLD, the Layer 1 and Layer 2 detections work together to prevent physical and logical unidirectional connections and the malfunctioning of other protocols.

A unidirectional link occurs whenever traffic sent by a local device is received by its neighbor but traffic from the neighbor is not received by the local device.

In normal mode, UDLD detects a unidirectional link when fiber strands in a fiber-optic port are misconnected and the Layer 1 mechanisms do not detect this misconnection. If the ports are connected correctly but the traffic is one way, UDLD does not detect the unidirectional link because the Layer 1 mechanism, which is supposed to detect this condition, does not do so. In this case, the logical link is considered undetermined, and UDLD does not disable the port.

When UDLD is in normal mode, if one of the fiber strands in a pair is disconnected, as long as autonegotiation is active, the link does not stay up because the Layer 1 mechanisms detects a physical problem with the link. In this case, UDLD does not take any action and the logical link is considered undetermined.

In aggressive mode, UDLD detects a unidirectional link by using the previous detection methods. UDLD in aggressive mode can also detect a unidirectional link on a point-to-point link on which no failure between the two devices is allowed. It can also detect a unidirectional link when one of these problems exists:

- On fiber-optic or twisted-pair links, one of the ports cannot send or receive traffic.
- On fiber-optic or twisted-pair links, one of the ports is down while the other is up.
- One of the fiber strands in the cable is disconnected.

In these cases, UDLD disables the affected port.

In a point-to-point link, UDLD hello packets can be considered as a heart beat whose presence guarantees the health of the link. Conversely, the loss of the heart beat means that the link must be shut down if it is not possible to reestablish a bidirectional link.

If both fiber strands in a cable are working normally from a Layer 1 perspective, UDLD in aggressive mode detects whether those fiber strands are connected correctly and whether traffic is flowing bidirectionally between the correct neighbors. This check cannot be performed by autonegotiation because autonegotiation operates at Layer 1.

Methods to Detect Unidirectional Links

UDLD operates by using two methods:

- Neighbor database maintenance

  UDLD learns about other UDLD-capable neighbors by periodically sending a hello packet (also called an advertisement or probe) on every active port to keep each device informed about its neighbors.

When the switch receives a hello message, it caches the information until the age time (hold time or time-to-live) expires. If the switch receives a new hello message before an older cache entry ages, the switch replaces the older entry with the new one.

Whenever a port is disabled and UDLD is running, whenever UDLD is disabled on a port, or whenever the switch is reset, UDLD clears all existing cache entries for the ports affected by the configuration change. UDLD sends at least one message to inform the neighbors to flush the part of their caches affected by the status change. The message is intended to keep the caches synchronized.

- Event-driven detection and echoing

  UDLD relies on echoing as its detection mechanism. Whenever a UDLD device learns about a new neighbor or receives a resynchronization request from an out-of-sync neighbor, it restarts the detection window on its side of the connection and sends echo messages in reply. Because this behavior is the same on all UDLD neighbors, the sender of the echoes expects to receive an echo in reply.

  If the detection window ends and no valid reply message is received, the link might shut down, depending on the UDLD mode. When UDLD is in normal mode, the link might be considered undetermined and might not be shut down. When UDLD is in aggressive mode, the link is considered unidirectional, and the port is disabled.

If UDLD in normal mode is in the advertisement or in the detection phase and all the neighbor cache entries are aged out, UDLD restarts the link-up sequence to resynchronize with any potentially out-of-sync neighbors.

If you enable aggressive mode when all the neighbors of a port have aged out either in the advertisement or in the detection phase, UDLD restarts the link-up sequence to resynchronize with any potentially out-of-sync neighbor. UDLD shuts down the port if, after the fast train of messages, the link state is still undetermined.

Figure 33-1 UDLD Detection of a Unidirectional Link

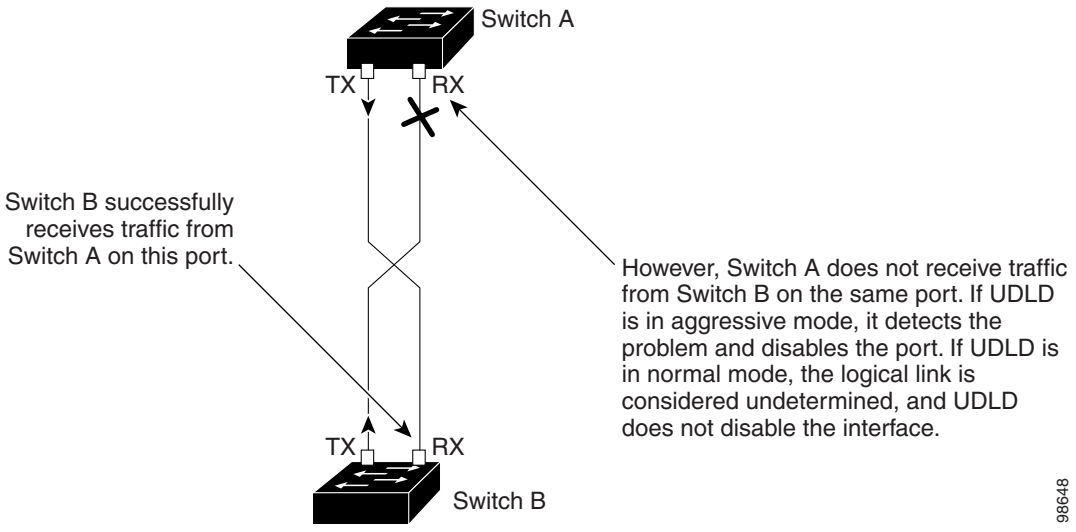

Default UDLD Settings

Table 33-1 Default UDLD Settings

Feature	Default Setting
UDLD global enable state	Globally disabled
UDLD per-port enable state for fiber-optic media	Disabled on all Ethernet fiber-optic ports
UDLD per-port enable state for twisted-pair (copper) media	Disabled on all Ethernet 10/100 and 1000BASE-TX ports
UDLD aggressive mode	Disabled

How to Configure UDLD

Enabling UDLD Globally

Follow these steps to enable UDLD in the aggressive or normal mode and to set the configurable message timer on all fiber-optic ports on the switch:

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**udld** {**aggressive** | **enable** | **message time** *message-timer-interval*}	Specifies the UDLD mode of operation: • **aggressive**—Enables UDLD in aggressive mode on all fiber-optic ports. • **enable**—Enables UDLD in normal mode on all fiber-optic ports on the switch. UDLD is disabled by default. An individual interface configuration overrides the setting of the **udld enable** global configuration command. For more information about aggressive and normal modes, see the "Modes of Operation" section on page 33-2. • **message time** *message-timer-interval*—Configures the period of time between UDLD probe messages on ports that are in the advertisement phase and are detected to be bidirectional. The range is from 1 to 90 seconds. **Note** This command affects fiber-optic ports only. Use the **udld** interface configuration command to enable UDLD on other port types. For more information, see the "Enabling UDLD on an Interface" section on page 33-5.
Step 3	**end**	Returns to privileged EXEC mode.

Enabling UDLD on an Interface

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the port to be enabled for UDLD, and enters interface configuration mode.
Step 3	**udld port** [**aggressive**]	UDLD is disabled by default. • **udld port**—Enables UDLD in normal mode on the specified port. • **udld port aggressive**—Enables UDLD in aggressive mode on the specified port. **Note** Use the **no udld port** interface configuration command to disable UDLD on a specified fiber-optic port. For more information about aggressive and normal modes, see the "Modes of Operation" section on page 33-2.
Step 4	**end**	Returns to privileged EXEC mode.

Setting and Resetting UDLD Parameters

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**udld reset**	(Optional) Resets all ports disabled by UDLD.
Step 3	**no udld** {**aggressive** | **enable**}	(Optional) Disables the UDLD ports.
Step 4	**udld** {**aggressive** | **enable**}	(Optional) Reenables the disabled ports.
Step 5	**errdisable recovery cause udld**	(Optional) Enables the timer to automatically recover from the UDLD error-disabled state.
Step 6	**errdisable recovery interval** *interval*	(Optional) Specifies the time to recover from the UDLD error-disabled state.
Step 7	**interface** *interface-id*	Enters interface configuration mode.
Step 8	**no udld port**	(Optional) Disables the UDLD fiber-optic port.
Step 9	**udld port** [**aggressive**]	(Optional) Re-enables the disabled fiber-optic port.
Step 10	**shutdown**	(Optional) Disables an interface port.
Step 11	**no shutdown**	(Optional) Restarts a disabled port.
Step 12	**show udld**	(Optional) Verifies your entries.

Maintaining and Monitoring UDLD

Command	Purpose
show udld [*interface-id*]	Displays UDLD status.

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 34

Configuring RMON

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for RMON

- You must configure SNMP on the switch to access RMON MIB objects.
- We recommend that you use a generic RMON console application on the network management station (NMS) to take advantage of the RMON network management capabilities.

Restrictions for RMON

- 64-bit counters are not supported for RMON alarms.

Information About RMON

RMON

RMON is an Internet Engineering Task Force (IETF) standard monitoring specification that allows various network agents and console systems to exchange network monitoring data. You can use the RMON feature with the Simple Network Management Protocol (SNMP) agent in the switch to monitor all the traffic flowing among switches on all connected LAN segments as shown in Figure 34-1.

Figure 34-1 Remote Monitoring Example

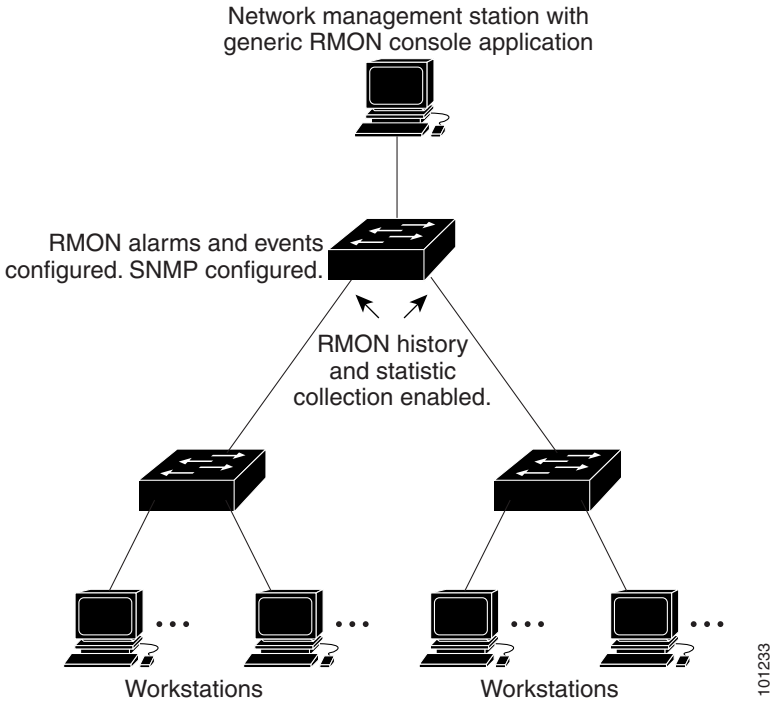

The switch supports these RMON groups (defined in RFC 1757):

- Statistics (RMON group 1)—Collects Ethernet statistics (including Fast Ethernet and Gigabit Ethernet statistics, depending on the switch type and supported interfaces) on an interface.
- History (RMON group 2)—Collects a history group of statistics on Ethernet ports (including Fast Ethernet and Gigabit Ethernet statistics, depending on the switch type and supported interfaces) for a specified polling interval.
- Alarm (RMON group 3)—Monitors a specific management information base (MIB) object for a specified interval, triggers an alarm at a specified value (rising threshold), and resets the alarm at another value (falling threshold). Alarms can be used with events; the alarm triggers an event, which can generate a log entry or an SNMP trap.
- Event (RMON group 9)—Specifies the action to take when an event is triggered by an alarm. The action can be to generate a log entry or an SNMP trap.

Because switches supported by this software release use hardware counters for RMON data processing, the monitoring is more efficient, and little processing power is required.

Note 64-bit counters are not supported for RMON alarms.

RMON is disabled by default; no alarms or events are configured.

How to Configure RMON

Configuring RMON Alarms and Events

You can configure your switch for RMON by using the command-line interface (CLI) or an SNMP-compatible network management station.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**rmon alarm** *number variable interval* {**absolute** | **delta**} **rising-threshold** *value* [*event-number*] **falling-threshold** *value* [*event-number*] [**owner** *string*]	Sets an alarm on a MIB object. • *number*—Specifies the alarm number. The range is 1 to 65535. • *variable*—Specifies the MIB object to monitor. • *interval*—Specifies the time in seconds the alarm monitors the MIB variable. The range is 1 to 4294967295 seconds. • Specifies the **absolute** keyword to test each MIB variable directly. Specifies the **delta** keyword to test the change between samples of a MIB variable. • *value*—Specifies a number at which the alarm is triggered and one for when the alarm is reset. The range for the rising threshold and falling threshold values is -2147483648 to 2147483647. • (Optional) *event-number*—Specifies the event number to trigger when the rising or falling threshold exceeds its limit. • (Optional) **owner** *string*—Specifies the owner of the alarm.
Step 3	**rmon event** *number* [**description** *string*] [**log**] [**owner** *string*] [**trap** *community*]	Adds an event in the RMON event table that is associated with an RMON event number. • *number*—Assigns an event number. The range is 1 to 65535. • (Optional) **description** *string*—Specifies a description of the event. • (Optional) **log**—Generates an RMON log entry when the event is triggered. • (Optional) **owner** *string*—Specifies the owner of this event. • (Optional) **trap** *community*—Enters the SNMP community string used for this trap.
Step 4	**end**	Returns to privileged EXEC mode.

Collecting Group History Statistics on an Interface

You must first configure RMON alarms and events to display collection information.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface on which to collect history, and enters interface configuration mode.
Step 3	**rmon collection history** *index* [**buckets** *bucket-number*] [**interval** *seconds*] [**owner** *ownername*]	Enables history collection for the specified number of buckets and time period. • *index*—Identifies the RMON group of statistics. The range is 1 to 65535. • (Optional) **buckets** *bucket-number*—Specifies the maximum number of buckets desired for the RMON collection history group of statistics. The range is 1 to 65535. The default is 50 buckets. • (Optional) **interval** *seconds*—Specifies the number of seconds in each polling cycle. The range is 1 to 3600. The default is 1800 seconds. • (Optional) **owner** *ownername*—Enters the name of the owner of the RMON group of statistics.
Step 4	**end**	Returns to privileged EXEC mode.

Collecting Group Ethernet Statistics on an Interface

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the interface on which to collect statistics, and enters interface configuration mode.
Step 3	**rmon collection stats** *index* [**owner** *ownername*]	Enables RMON statistic collection on the interface. • *index*—Specifies the RMON group of statistics. The range is from 1 to 65535. • (Optional) **owner** *ownername*—Enters the name of the owner of the RMON group of statistics.
Step 4	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining RMON

Table 34-1 Commands for Displaying RMON Status

Command	Purpose
show rmon	Displays general RMON statistics.
show rmon alarms	Displays the RMON alarm table.
show rmon events	Displays the RMON event table.
show rmon history	Displays the RMON history table.
show rmon statistics	Displays the RMON statistics table.

Configuration Examples for RMON

Configuring an RMON Alarm Number: Example

The following example shows how to configure an RMON alarm number:

```
Switch(config)# rmon alarm 10 ifEntry.20.1 20 delta rising-threshold 15 1
falling-threshold 0 owner jjohnson
```

The alarm monitors the MIB variable *ifEntry.20.1* once every 20 seconds until the alarm is disabled and checks the change in the variable's rise or fall. If the *ifEntry.20.1* value shows a MIB counter increase of 15 or more, such as from 100000 to 100015, the alarm is triggered. The alarm in turn triggers event number 1, which is configured with the **rmon event** command. Possible events can include a log entry or an SNMP trap. If the *ifEntry.20.1* value changes by 0, the alarm is reset and can be triggered again.

Creating an RMON Event Number: Example

The following example creates RMON event number 1:

```
Switch(config)# rmon event 1 log trap eventtrap description "High ifOutErrors" owner
jjones
```

The event is defined as *High ifOutErrors* and generates a log entry when the event is triggered by the alarm. The user *jjones* owns the row that is created in the event table by this command. This example also generates an SNMP trap when the event is triggered.

Configuring RMON Statistics: Example

This example shows how to collect RMON statistics for the owner *root*:

```
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# rmon collection stats 2 owner root
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands Cisco IOS system management commands	*Cisco IOS Configuration Fundamentals Command Reference*
SNMP configuration	Chapter 36, "Configuring SNMP"
Alarm and event interaction	RFC 1757

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 35

Configuring System Message Logging

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Restrictions for System Message Logging

- Logging messages to the console at a high rate can result in high CPU utilization and adversely affect how the switch operates.

Information About System Message Logging

System Message Logging

By default, a switch sends the output from system messages and **debug** privileged EXEC commands to a logging process. The logging process controls the distribution of logging messages to various destinations, such as the logging buffer, terminal lines, or a UNIX syslog server, depending on your configuration. The process also sends messages to the console.

Note The syslog format is compatible with 4.3 BSD UNIX.

When the logging process is disabled, messages are sent only to the console. The messages are sent as they are generated, so message and debug output are interspersed with prompts or output from other commands. Messages appear on the console after the process that generated them has finished.

You can set the severity level of the messages to control the type of messages displayed on the consoles and each of the destinations. You can time-stamp log messages or set the syslog source address to enhance real-time debugging and management. For information on possible messages, see the system message guide for this release.

You can access logged system messages by using the switch command-line interface (CLI) or by saving them to a properly configured syslog server. The switch software saves syslog messages in an internal buffer.

You can remotely monitor system messages by viewing the logs on a syslog server or by accessing the switch through Telnet or through the console port.

System Log Message Format

System log messages can contain up to 80 characters and a percent sign (%), which follows the optional sequence number or time-stamp information, if configured. Messages appear in this format:

seq no:timestamp: %facility-severity-MNEMONIC:description

The part of the message preceding the percent sign depends on the setting of the **service sequence-numbers**, **service timestamps log datetime**, **service timestamps log datetime** [**localtime**] [**msec**] [**show-timezone**], or **service timestamps log uptime** global configuration command.

Table 35-1 System Log Message Elements

Element	Description
seq no:	Stamps log messages with a sequence number only if the **service sequence-numbers** global configuration command is configured. For more information, see the "Enabling and Disabling Sequence Numbers in Log Messages" section on page 35-8.
timestamp formats: *mm/dd hh:mm:ss* or *hh:mm:ss* (short uptime) or *d h* (long uptime)	Date and time of the message or event. This information appears only if the **service timestamps log** [**datetime** | **log**] global configuration command is configured. For more information, see the "Enabling and Disabling Time Stamps on Log Messages" section on page 35-8.
facility	The facility to which the message refers (for example, SNMP, SYS, and so forth). For a list of supported facilities, see Table 35-3 on page 35-4.
severity	Single-digit code from 0 to 7 that is the severity of the message. For a description of the severity levels, see Table 35-2 on page 35-3.
MNEMONIC	Text string that uniquely describes the message.
description	Text string containing detailed information about the event being reported.

Log Messages

You can synchronize unsolicited messages and **debug** privileged EXEC command output with solicited device output and prompts for a specific console port line or virtual terminal line. You can identify the types of messages to be output asynchronously based on the level of severity. You can also configure the maximum number of buffers for storing asynchronous messages for the terminal after which messages are dropped.

When synchronous logging of unsolicited messages and **debug** command output is enabled, unsolicited device output appears on the console or printed after solicited device output appears or is printed. Unsolicited messages and **debug** command output appears on the console after the prompt for user input is returned. Therefore, unsolicited messages and **debug** command output are not interspersed with solicited device output and prompts. After the unsolicited messages appear, the console again displays the user prompt.

Message Severity Levels

Note Specifying a *level* causes messages at that level and numerically lower levels to appear at the destination.

To disable logging to the console, use the **no logging console** global configuration command. To disable logging to a terminal other than the console, use the **no logging monitor** global configuration command. To disable logging to syslog servers, use the **no logging trap** global configuration command.

Table 35-2 describes the *level* keywords. It also lists the corresponding UNIX syslog definitions from the most severe level to the least severe level.

Table 35-2 Message Logging Level Keywords

Level Keyword	Level	Description	Syslog Definition
emergencies	0	System unstable	LOG_EMERG
alerts	1	Immediate action needed	LOG_ALERT
critical	2	Critical conditions	LOG_CRIT
errors	3	Error conditions	LOG_ERR
warnings	4	Warning conditions	LOG_WARNING
notifications	5	Normal but significant condition	LOG_NOTICE
informational	6	Informational messages only	LOG_INFO
debugging	7	Debugging messages	LOG_DEBUG

The software generates these categories of messages:

- Error messages about software or hardware malfunctions, displayed at levels **warnings** through **emergencies**. These types of messages mean that the functionality of the switch is affected. For information on how to recover from these malfunctions, see the system message guide for this release.
- Output from the **debug** commands, displayed at the **debugging** level. Debug commands are typically used only by the Technical Assistance Center.
- Interface up or down transitions and system restart messages, displayed at the **notifications** level. This message is only for information; switch functionality is not affected.

Configuring UNIX Syslog Servers

The next sections describe how to configure the UNIX server syslog daemon and how to define the UNIX system logging facility.

Logging Messages to a UNIX Syslog Daemon

Before you can send system log messages to a UNIX syslog server, you must configure the syslog daemon on a UNIX server. This procedure is optional.

Note Some recent versions of UNIX syslog daemons no longer accept by default syslog packets from the network. If this is the case with your system, use the UNIX **man syslogd** command to decide what options must be added to or removed from the syslog command line to enable logging of remote syslog messages.

Log in as root, and perform these steps:

Step 1 Add a line such as the following to the file /etc/syslog.conf:

```
local7.debug /usr/adm/logs/cisco.log
```

The **local7** keyword specifies the logging facility to be used; see Table 35-3 on page 35-4 for information on the facilities. The **debug** keyword specifies the syslog level; see Table 35-2 on page 35-3 for information on the severity levels. The syslog daemon sends messages at this level or at a more severe level to the file specified in the next field. The file must already exist, and the syslog daemon must have permission to write to it.

Step 2 Create the log file by entering these commands at the UNIX shell prompt:

```
$ touch /var/log/cisco.log
$ chmod 666 /var/log/cisco.log
```

Step 3 Make sure the syslog daemon reads the new changes:

```
$ kill -HUP `cat /etc/syslog.pid`
```

For more information, see the **man syslog.conf** and **man syslogd** commands on your UNIX system.

Table 35-3 lists the UNIX system facilities supported by the software. For more information about these facilities, consult the operator's manual for your UNIX operating system.

Table 35-3 Logging Facility-Type Keywords

Facility Type Keyword	Description
auth	Authorization system
cron	Cron facility
daemon	System daemon
kern	Kernel
local0-7	Locally defined messages
lpr	Line printer system
mail	Mail system
news	USENET news
sys9-14	System use
syslog	System log

Table 35-3 Logging Facility-Type Keywords (continued)

Facility Type Keyword	Description
user	User process
uucp	UNIX-to-UNIX copy system

Default System Message Logging Configuration

Table 35-4 Default System Message Logging Configuration

Feature	Default Setting
System message logging to the console	Enabled.
Console severity	Debugging (and numerically lower levels; see Table 35-2 on page 35-3).
Logging file configuration	No filename specified.
Logging buffer size	4096 bytes.
Logging history size	1 message.
Time stamps	Disabled.
Synchronous logging	Disabled.
Logging server	Disabled.
Syslog server IP address	None configured.
Configuration change logger	Disabled.
Server facility	Local7 (see Table 35-3 on page 35-4).
Server severity	Informational (and numerically lower levels; see Table 35-2 on page 35-3).

How to Configure System Message Logging

Disabling Message Logging

Message logging is enabled by default. It must be enabled to send messages to any destination other than the console. When enabled, log messages are sent to a logging process, which logs messages to designated locations asynchronously to the processes that generated the messages.

Disabling the logging process can slow down the switch because a process must wait until the messages are written to the console before continuing. When the logging process is disabled, messages appear on the console as soon as they are produced, often appearing in the middle of command output.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.

	Command	Purpose
Step 2	**no logging console**	Disables message logging.
Step 3	**end**	Returns to privileged EXEC mode.

Setting the Message Display Destination Device

If message logging is enabled, you can send messages to specific locations in addition to the console. Beginning in privileged EXEC mode, use one or more of the following commands to specify the locations that receive messages:

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**logging buffered** [*size*]	Logs messages to an internal buffer on the switch. The range is 4096 to 2147483647 bytes. The default buffer size is 4096 bytes. If the switch fails, the log file is lost unless you had previously saved it to flash memory. See Step 4. **Note** Do not make the buffer size too large because the switch could run out of memory for other tasks. Use the **show memory** privileged EXEC command to view the free processor memory on the switch. However, this value is the maximum available, and the buffer size should *not* be set to this amount.
Step 3	**logging** *host*	Logs messages to a UNIX syslog server host. *host*—Specifies the name or IP address of the host to be used as the syslog server. To build a list of syslog servers that receive logging messages, enter this command more than once.
Step 4	**logging file flash:***filename* [*max-file-size* [*min-file-size*]] [*severity-level-number* | *type*]	Stores log messages in a file in flash memory. • *filename*—Enters the log message filename. • (Optional) *max-file-size*—Specifies the maximum logging file size. The range is 4096 to 2147483647. The default is 4096 bytes. • (Optional) *min-file-size*—Specifies the minimum logging file size. The range is 1024 to 2147483647. The default is 2048 bytes. • (Optional) *severity-level-number* | *type*—Specifies either the logging severity level or the logging type. The severity range is 0 to 7. For a list of logging type keywords, see Table 35-2 on page 35-3. By default, the log file receives debugging messages and numerically lower levels.
Step 5	**end**	Returns to privileged EXEC mode.
Step 6	**terminal monitor**	Logs messages to a nonconsole terminal during the current session. Terminal parameter-setting commands are set locally and do not remain in effect after the session has ended. You must perform this step for each session to see the debugging messages.

Synchronizing Log Messages

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**line** [**console** | **vty**] *line-number* [*ending-line-number*]	Specifies the line to be configured for synchronous logging of messages. • Use the **console** keyword for configurations that occur through the switch console port. • Use the **line vty** *line-number* command to specify which vty lines are to have synchronous logging enabled. You use a vty connection for configurations that occur through a Telnet session. The range of line numbers is from 0 to 15. You can change the setting of all 16 vty lines at once by entering: **line vty 0 15** Or you can change the setting of the single vty line being used for your current connection. For example, to change the setting for vty line 2, enter: **line vty 2** When you enter this command, the mode changes to line configuration.
Step 3	**logging synchronous** [**level** [*severity-level* | **all**] | **limit** *number-of-buffers*]	Enables synchronous logging of messages. • (Optional) **level** *severity-level*—Specifies the message severity level. Messages with a severity level equal to or higher than this value are printed asynchronously. Low numbers mean greater severity and high numbers mean lesser severity. The default is 2. • (Optional) **level all**—Specifies that all messages are printed asynchronously regardless of the severity level. • (Optional) **limit** *number-of-buffers*—Specifies the number of buffers to be queued for the terminal after which new messages are dropped. The range is 0 to 2147483647. The default is 20.
Step 4	**end**	Returns to privileged EXEC mode.

Enabling and Disabling Time Stamps on Log Messages

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**service timestamps log uptime** or **service timestamps log datetime** [**msec**] [**localtime**] [**show-timezone**]	Enables log time stamps. The first command enables time stamps on log messages, showing the time since the system was rebooted. The second command enables time stamps on log messages. Depending on the options selected, the time stamp can include the date, time in milliseconds relative to the local time-zone, and the time zone name.
Step 3	**end**	Returns to privileged EXEC mode.

Enabling and Disabling Sequence Numbers in Log Messages

Because there is a chance that more than one log message can have the same time stamp, you can display messages with sequence numbers so that you can unambiguously see a single message. By default, sequence numbers in log messages are not displayed.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**service sequence-numbers**	Enables sequence numbers.
Step 3	**end**	Returns to privileged EXEC mode.

Defining the Message Severity Level

You can limit messages displayed to the selected device by specifying the severity level of the message, which are described in Table 35-2.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**logging console** *level*	Limits messages logged to the console. By default, the console receives debugging messages and numerically lower levels.
Step 3	**logging monitor** *level*	Limits messages logged to the terminal lines. By default, the terminal receives debugging messages and numerically lower levels.
Step 4	**logging trap** *level*	Limits messages logged to the syslog servers. By default, syslog servers receive informational messages and numerically lower levels.
Step 5	**end**	Returns to privileged EXEC mode.

Limiting Syslog Messages Sent to the History Table and to SNMP

If you enabled syslog message traps to be sent to an SNMP network management station by using the **snmp-server enable trap** global configuration command, you can change the level of messages sent and stored in the switch history table. You also can change the number of messages that are stored in the history table.

Messages are stored in the history table because SNMP traps are not guaranteed to reach their destination. By default, one message of the level **warning** and numerically lower levels (see Table 35-2 on page 35-3) are stored in the history table even if syslog traps are not enabled.

When the history table is full (it contains the maximum number of message entries specified with the **logging history size** global configuration command), the oldest message entry is deleted from the table to allow the new message entry to be stored.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**logging history** *level*	Changes the default level of syslog messages stored in the history file and sent to the SNMP server. By default, **warnings**, **errors**, **critical**, **alerts**, and **emergencies** messages are sent.
Step 3	**logging history size** *number*	Specifies the number of syslog messages that can be stored in the history table. The default is to store one message. The range is 0 to 500 messages.
Step 4	**end**	Returns to privileged EXEC mode.

Enabling the Configuration-Change Logger

You can enable a configuration logger to keep track of configuration changes made with the command-line interface (CLI). When you enter the **logging enable** configuration-change logger configuration command, the log records the session, the user, and the command that was entered to change the configuration. You can configure the size of the configuration log from 1 to 1000 entries (the default is 100).

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**archive**	Enters archive configuration mode.
Step 3	**log config**	Enters configuration-change logger configuration mode.
Step 4	**logging enable**	Enables configuration change logging.
Step 5	**logging size** *entries*	(Optional) Configures the number of entries retained in the configuration log. The range is from 1 to 1000. The default is 100. **Note** When the configuration log is full, the oldest log entry is removed each time a new entry is entered.
Step 6	**end**	Returns to privileged EXEC mode.

Configuring the UNIX System Logging Facility

When sending system log messages to an external device, you can cause the switch to identify its messages as originating from any of the UNIX syslog facilities.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**logging** *host*	Logs messages to a UNIX syslog server host by entering its IP address. To build a list of syslog servers that receive logging messages, enter this command more than once.
Step 3	**logging trap** *level*	Limits messages logged to the syslog servers. Be default, syslog servers receive informational messages and lower. See Table 35-2 on page 35-3 for *level* keywords.
Step 4	**logging facility** *facility-type*	Configures the syslog facility. See Table 35-3 on page 35-4 for *facility-type* keywords. The default is **local7**.
Step 5	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining the System Message Log

Command	Purpose
show logging	Displays logging messages.
show archive log config	Displays the configuration log.

Configuration Examples for the System Message Log

System Message: Example

This example shows a partial switch system message:

```
00:00:46: %LINK-3-UPDOWN: Interface Port-channel1, changed state to up
00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/1, changed state to up
00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/2, changed state to up
00:00:48: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1, changed state to down
00:00:48: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/1, changed
state to down 2
*Mar  1 18:46:11: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)
18:47:02: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)
*Mar  1 18:48:50.483 UTC: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)
```

Logging Display: Examples

This example shows part of a logging display with the **service timestamps log datetime** global configuration command enabled:

```
*Mar  1 18:46:11: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)
```

This example shows part of a logging display with the s**ervice timestamps log uptime** global configuration command enabled:

```
00:00:46: %LINK-3-UPDOWN: Interface Port-channel1, changed state to up
```

This example shows part of a logging display with sequence numbers enabled:

```
000019: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)
```

Enabling the Logger: Example

This example shows how to enable the configuration-change logger and to set the number of entries in the log to 500.

```
Switch(config)# archive
Switch(config-archive)# log config
Switch(config-archive-log-cfg)# logging enable
Switch(config-archive-log-cfg)# logging size 500
Switch(config-archive-log-cfg)# end
```

Configuration Log Output: Example

This is an example of output for the configuration log:

```
Switch# show archive log config all
 idx   sess           user@line      Logged command
   38    11   unknown user@vty3     |no aaa authorization config-commands
   39    12   unknown user@vty3     |no aaa authorization network default group radius
   40    12   unknown user@vty3     |no aaa accounting dot1x default start-stop group
radius
   41    13   unknown user@vty3     |no aaa accounting system default
   42    14           temi@vty4     |interface GigabitEthernet4/0/1
   43    14           temi@vty4     | switchport mode trunk
   44    14           temi@vty4     | exit
   45    16           temi@vty5     |interface FastEthernet5/0/1
   46    16           temi@vty5     | switchport mode trunk
   47    16           temi@vty5     | exit
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands Cisco IOS system management commands	*Cisco IOS Configuration Fundamentals Command Reference*
Syslog server configuration steps	"Configuring the UNIX System Logging Facility" section on page 35-10

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 36

Configuring SNMP

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for SNMP

An SNMP *group* is a table that maps SNMP users to SNMP views. An SNMP *user* is a member of an SNMP group. An SNMP *host* is the recipient of an SNMP trap operation. An SNMP *engine ID* is a name for the local or remote SNMP engine.

- If the switch starts and the switch startup configuration has at least one **snmp-server** global configuration command, the SNMP agent is enabled.
- When configuring an SNMP group, do not specify a notify view. The **snmp-server host** global configuration command autogenerates a notify view for the user and then adds it to the group associated with that user. Modifying the group's notify view affects all users associated with that group. See the *Cisco IOS Network Management Command Reference* for information about when you should configure notify views.
- To configure a remote user, specify the IP address or port number for the remote SNMP agent of the device where the user resides.
- Before you configure remote users for a particular agent, configure the SNMP engine ID, using the **snmp-server engineID** global configuration with the **remote** option. The remote agent's SNMP engine ID and user password are used to compute the authentication and privacy digests. If you do not configure the remote engine ID first, the configuration command fails.

Restrictions for SNMP

- When configuring SNMP informs, you need to configure the SNMP engine ID for the remote agent in the SNMP database before you can send proxy requests or informs to it.
- If a local user is not associated with a remote host, the switch does not send informs for the **auth** (authNoPriv) and the **priv** (authPriv) authentication levels.

- Changing the value of the SNMP engine ID has important implications. A user's password (entered on the command line) is converted to an MD5 or SHA security digest based on the password and the local engine ID. The command-line password is then destroyed, as required by RFC 2274. Because of this deletion, if the value of the engine ID changes, the security digests of SNMPv3 users become invalid, and you need to reconfigure SNMP users by using the **snmp-server user** *username* global configuration command. Similar restrictions require the reconfiguration of community strings when the engine ID changes.

Information About SNMP

SNMP

The Simple Network Management Protocol (SNMP) is an application-layer protocol that provides a message format for communication between managers and agents. The SNMP system consists of an SNMP manager, an SNMP agent, and a MIB. The SNMP manager can be part of a network management system (NMS) such as CiscoWorks. The agent and MIB reside on the switch. To configure SNMP on the switch, you define the relationship between the manager and the agent.

The SNMP agent contains MIB variables whose values the SNMP manager can request or change. A manager can get a value from an agent or store a value into the agent. The agent gathers data from the MIB, the repository for information about device parameters and network data. The agent can also respond to a manager's requests to get or set data.

An agent can send unsolicited traps to the manager. Traps are messages alerting the SNMP manager to a condition on the network. Traps can mean improper user authentication, restarts, link status (up or down), MAC address tracking, closing of a TCP connection, loss of connection to a neighbor, or other significant events.

SNMP Versions

This software release supports these SNMP versions:

- SNMPv1—The Simple Network Management Protocol, a Full Internet Standard, defined in RFC 1157.
- SNMPv2C replaces the Party-based Administrative and Security Framework of SNMPv2Classic with the community-string-based Administrative Framework of SNMPv2C while retaining the bulk retrieval and improved error handling of SNMPv2Classic. It has these features:
  - SNMPv2—Version 2 of the Simple Network Management Protocol, a Draft Internet Standard, defined in RFCs 1902 through 1907.
  - SNMPv2C—The community-string-based Administrative Framework for SNMPv2, an Experimental Internet Protocol defined in RFC 1901.
- SNMPv3—Version 3 of the SNMP is an interoperable standards-based protocol defined in RFCs 2273 to 2275. SNMPv3 provides secure access to devices by authenticating and encrypting packets over the network and includes these security features:
  - Message integrity—Ensures that a packet was not tampered with in transit.
  - Authentication—Determines that the message is from a valid source.

– Encryption—Mixes the contents of a package to prevent it from being read by an unauthorized source.

Note To select encryption, enter the **priv** keyword. This keyword is available only when the cryptographic (encrypted) software image is installed.

Both SNMPv1 and SNMPv2C use a community-based form of security. The community of managers able to access the agent's MIB is defined by an IP address access control list and password.

SNMPv2C includes a bulk retrieval mechanism and more detailed error message reporting to management stations. The bulk retrieval mechanism retrieves tables and large quantities of information, minimizing the number of round-trips required. The SNMPv2C improved error-handling includes expanded error codes that distinguish different kinds of error conditions; these conditions are reported through a single error code in SNMPv1. Error return codes in SNMPv2C report the error type.

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy set up for a user and the group within which the user resides. A security level is the permitted level of security within a security model. A combination of the security level and the security model determine which security mechanism is used when handling an SNMP packet. Available security models are SNMPv1, SNMPv2C, and SNMPv3.

Table 36-1 identifies the characteristics of the different combinations of security models and levels.

Table 36-1 SNMP Security Models and Levels

Model	Level	Authentication	Encryption	Result
SNMPv1	noAuthNoPriv	Community string	No	Uses a community string match for authentication.
SNMPv2C	noAuthNoPriv	Community string	No	Uses a community string match for authentication.
SNMPv3	noAuthNoPriv (requires the LAN Base image)	Username	No	Uses a username match for authentication.
SNMPv3	authNoPriv (requires the LAN Base image)	Message Digest 5 (MD5) or Secure Hash Algorithm (SHA)	No	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms.
SNMPv3	authPriv (requires the LAN Base image)	MD5 or SHA	Data Encryption Standard (DES) or Advanced Encryption Standard (AES)	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms. Allows specifying the User-based Security Model (USM) with these encryption algorithms: • DES 56-bit encryption in addition to authentication based on the CBC-DES (DES-56) standard. • 3DES 168-bit encryption • AES 128-bit, 192-bit, or 256-bit encryption

You must configure the SNMP agent to use the SNMP version supported by the management station. Because an agent can communicate with multiple managers, you can configure the software to support communications using SNMPv1, SNMPv2C, or SNMPv3.

SNMP Manager Functions

The SNMP manager uses information in the MIB to perform the operations described in Table 36-2.

Table 36-2 SNMP Operations

Operation	Description
get-request	Retrieves a value from a specific variable.
get-next-request	Retrieves a value from a variable within a table.[1]
get-bulk-request[2]	Retrieves large blocks of data, such as multiple rows in a table, that would otherwise require the transmission of many small blocks of data.
get-response	Replies to a get-request, get-next-request, and set-request sent by an NMS.
set-request	Stores a value in a specific variable.
trap	An unsolicited message sent by an SNMP agent to an SNMP manager when some event has occurred.

1. With this operation, an SNMP manager does not need to know the exact variable name. A sequential search is performed to find the needed variable from within a table.
2. The **get-bulk** command only works with SNMPv2 or later.

SNMP Agent Functions

The SNMP agent responds to SNMP manager requests as follows:

- Get a MIB variable—The SNMP agent begins this function in response to a request from the NMS. The agent retrieves the value of the requested MIB variable and responds to the NMS with that value.
- Set a MIB variable—The SNMP agent begins this function in response to a message from the NMS. The SNMP agent changes the value of the MIB variable to the value requested by the NMS.

The SNMP agent also sends unsolicited trap messages to notify an NMS that a significant event has occurred on the agent. Examples of trap conditions include, but are not limited to, when a port or module goes up or down, when spanning-tree topology changes occur, and when authentication failures occur.

SNMP Community Strings

SNMP community strings authenticate access to MIB objects and function as embedded passwords. In order for the NMS to access the switch, the community string definitions on the NMS must match at least one of the three community string definitions on the switch.

A community string can have one of these attributes:

- Read-only (RO)—Gives read access to authorized management stations to all objects in the MIB except the community strings, but does not allow write access.
- Read-write (RW)—Gives read and write access to authorized management stations to all objects in the MIB, but does not allow access to the community strings.

When a cluster is created, the command switch manages the exchange of messages among member switches and the SNMP application. The Network Assistant software appends the member switch number (*@esN*, where *N* is the switch number) to the first configured RW and RO community strings on the command switch and propagates them to the member switches. For more information, see Chapter 6, "Configuring Switch Clusters" and see *Getting Started with Cisco Network Assistant*, available on Cisco.com.

Using SNMP to Access MIB Variables

An example of an NMS is the CiscoWorks network management software. CiscoWorks 2000 software uses the switch MIB variables to set device variables and to poll devices on the network for specific information. The results of a poll can be displayed as a graph and analyzed to troubleshoot internetworking problems, increase network performance, verify the configuration of devices, monitor traffic loads, and more.

As shown in Figure 36-1, the SNMP agent gathers data from the MIB. The agent can send traps, or notification of certain events, to the SNMP manager, which receives and processes the traps. Traps alert the SNMP manager to a condition on the network such as improper user authentication, restarts, link status (up or down), MAC address tracking, and so forth. The SNMP agent also responds to MIB-related queries sent by the SNMP manager in *get-request*, *get-next-request*, and *set-request* format.

Figure 36-1 SNMP Network

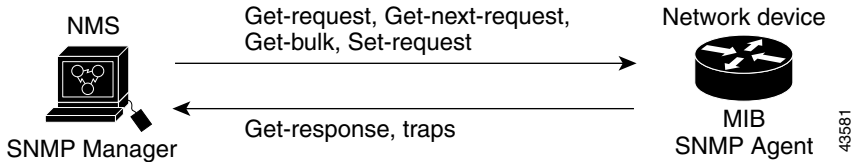

SNMP Notifications

SNMP allows the switch to send notifications to SNMP managers when particular events occur. SNMP notifications can be sent as traps or inform requests. In command syntax, unless there is an option in the command to select either traps or informs, the keyword **traps** refers to either traps or informs, or both. Use the **snmp-server host** command to specify whether to send SNMP notifications as traps or informs.

Note SNMPv1 does not support informs.

Traps are unreliable because the receiver does not send an acknowledgment when it receives a trap, and the sender cannot determine if the trap was received. When an SNMP manager receives an inform request, it acknowledges the message with an SNMP response protocol data unit (PDU). If the sender does not receive a response, the inform request can be sent again. Because they can be resent, informs are more likely than traps to reach their intended destination.

The characteristics that make informs more reliable than traps also consume more resources in the switch and in the network. Unlike a trap, which is discarded as soon as it is sent, an inform request is held in memory until a response is received or the request times out. Traps are sent only once, but an inform might be resent or retried several times. The retries increase traffic and contribute to a higher overhead on the network. Therefore, traps and informs require a trade-off between reliability and resources. If it is important that the SNMP manager receive every notification, use inform requests. If traffic on the network or memory in the switch is a concern and notification is not required, use traps.

SNMP ifIndex MIB Object Values

In an NMS, the IF-MIB generates and assigns an interface index (ifIndex) object value that is a unique number greater than zero to identify a physical or a logical interface. When the switch reboots or the switch software is upgraded, the switch uses this same value for the interface. For example, if the switch assigns a port 2 an ifIndex value of 10003, this value is the same after the switch reboots.

The switch uses one of the values in Table 36-3 to assign an ifIndex value to an interface.

Table 36-3 ifIndex Values

Interface Type	ifIndex Range
SVI	1–4999
EtherChannel	5001–5048
Physical (such as Gigabit Ethernet or SFP-module interfaces) based on type and port numbers	10000–14500
Null	10501
Loopback and Tunnel	24567 +

Note The switch might not use sequential values within a range.

Community Strings

You use the SNMP community string to define the relationship between the SNMP manager and the agent. The community string acts like a password to permit access to the agent on the switch. Optionally, you can specify one or more of these characteristics associated with the string:

- An access list of IP addresses of the SNMP managers that are permitted to use the community string to gain access to the agent
- A MIB view, which defines the subset of all MIB objects accessible to the given community
- Read and write or read-only permission for the MIB objects accessible to the community

SNMP Notifications

A trap manager is a management station that receives and processes traps. Traps are system alerts that the switch generates when certain events occur. By default, no trap manager is defined, and no traps are sent. Switches running this Cisco IOS release can have an unlimited number of trap managers.

Note Many commands use the word *traps* in the command syntax. Unless there is an option in the command to select either traps or informs, the keyword **traps** refers to traps, informs, or both. Use the **snmp-server host** global configuration command to specify whether to send SNMP notifications as traps or informs.

This table describes the supported switch traps (notification types). You can enable any or all of these traps and configure a trap manager to receive them. To enable the sending of SNMP inform notifications, use the **snmp-server enable traps** global configuration command combined with the **snmp-server host** *host-addr* **informs** global configuration command.

Table 36-4 Switch Notification Types

Notification Type Keyword	Description
bridge	Generates STP bridge MIB traps.
config	Generates a trap for SNMP configuration changes.
copy-config	Generates a trap for SNMP copy configuration changes.
entity	Generates a trap for SNMP entity changes.
cpu threshold	Allows CPU-related traps.
envmon	Generates environmental monitor traps. You can enable any or all of these environmental traps: fan, shutdown, status, supply, temperature.
errdisable	Generates a trap for an error-disabled VLAN port. You can also set a maximum trap rate per minute. The range is from 0 to 10000; the default is 0, which means there is no rate limit.
flash	Generates SNMP FLASH notifications.
hsrp	Generates a trap for Hot Standby Router Protocol (HSRP) changes.
ipmulticast	Generates a trap for IP multicast routing changes.
mac-notification	Generates a trap for MAC address notifications.
msdp	Generates a trap for Multicast Source Discovery Protocol (MSDP) changes.
ospf	Generates a trap for Open Shortest Path First (OSPF) changes. You can enable any or all of these traps: Cisco specific, errors, link-state advertisement, rate limit, retransmit, and state changes.
pim	Generates a trap for Protocol-Independent Multicast (PIM) changes. You can enable any or all of these traps: invalid PIM messages, neighbor changes, and rendezvous point (RP)-mapping changes.
port-security	Generates SNMP port security traps. You can also set a maximum trap rate per second. The range is from 0 to 1000; the default is 0, which means that there is no rate limit. **Note** When you configure a trap by using the notification type **port-security**, configure the port security trap first, and then configure the port security trap rate: • **snmp-server enable traps port-security** • **snmp-server enable traps port-security trap-rate** *rate*
rtr	Generates a trap for the SNMP Response Time Reporter (RTR).
snmp	Generates a trap for SNMP-type notifications for authentication, cold start, warm start, link up or link down.
storm-control	Generates a trap for SNMP storm control. You can also set a maximum trap rate per minute. The range is from 0 to 1000; the default is 0 (no limit is imposed; a trap is sent at every occurrence).
stpx	Generates SNMP STP Extended MIB traps.
syslog	Generates SNMP syslog traps.
tty	Generates a trap for TCP connections. This trap is enabled by default.
vlan-membership	Generates a trap for SNMP VLAN membership changes.
vlancreate	Generates SNMP VLAN created traps.

Table 36-4 Switch Notification Types (continued)

Notification Type Keyword	Description
vlandelete	Generates SNMP VLAN deleted traps.
vtp	Generates a trap for VLAN Trunking Protocol (VTP) changes.

Note Though visible in the command-line help strings, the **fru-ctrl, insertion**, and **removal** keywords are not supported.

You can use the **snmp-server host** global configuration command to a specific host to receive the notification types listed in Table 36-4.

Default SNMP Settings

Table 36-5 Default SNMP Settings

Feature	Default Setting
SNMP agent	Disabled[1].
SNMP trap receiver	None configured.
SNMP traps	None enabled except the trap for TCP connections (**tty**).
SNMP version	If no **version** keyword is present, the default is Version 1.
SNMPv3 authentication	If no keyword is entered, the default is the **noauth** (noAuthNoPriv) security level.
SNMP notification type	If no type is specified, all notifications are sent.

1. This is the default when the switch starts and the startup configuration does not have any **snmp-server** global configuration commands.

How to Configure SNMP

Disabling the SNMP Agent

The **no snmp-server** global configuration command disables all running versions (Version 1, Version 2C, and Version 3) on the device. No specific Cisco IOS command exists to enable SNMP. The first **snmp-server** global configuration command that you enter enables all versions of SNMP.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**no snmp-server**	Disables the SNMP agent operation.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring Community Strings

Note To disable access for an SNMP community, set the community string for that community to the null string (do not enter a value for the community string).

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**snmp-server community** *string* [**view** *view-name*] [**ro** | **rw**] [*access-list-number*]	Configures the community string. **Note** The @ symbol is used for delimiting the context information. Avoid using the @ symbol as part of the SNMP community string when configuring this command. • *string*—Specifies a string that acts like a password and permits access to the SNMP protocol. You can configure one or more community strings of any length. • (Optional) **view**—Specifies the view record accessible to the community. • (Optional) Specifies either read-only (**ro**) if you want authorized management stations to retrieve MIB objects, or specifies read-write (**rw**) if you want authorized management stations to retrieve and modify MIB objects. By default, the community string permits read-only access to all objects. • (Optional) *access-list-number*—Specifies an IP standard access list numbered from 1 to 99 and 1300 to 1999.
Step 3	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*]	(Optional) If you specified an IP standard access list number in Step 2, then create the list, repeating the command as many times as necessary. • *access-list-number*—Specifies the access list number specified in Step 2. • **deny** — Denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • *source*—Specifies the IP address of the SNMP managers that are permitted to use the community string to gain access to the agent. • (Optional) *source-wildcard*—Specifies the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. The access list is always terminated by an implicit deny statement for everything.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring SNMP Groups and Users

You can specify an identification name (engine ID) for the local or remote SNMP server engine on the switch. You can configure an SNMP server group that maps SNMP users to SNMP views, and you can add new users to the SNMP group.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**snmp-server engineID** {**local** *engineid-string* | **remote** *ip-address* [**udp-port** *port-number*] *engineid-string*}	Configures a name for either the local or remote copy of SNMP. • The *engineid-string* is a 24-character ID string with the name of the copy of SNMP. You need not specify the entire 24-character engine ID if it has trailing zeros. Specify only the portion of the engine ID up to the point where only zeros remain in the value. For example, to configure an engine ID of 123400000000000000000000, you can enter this: **snmp-server engineID local 1234** • If you select **remote**, specify the *ip-address* of the device that contains the remote copy of SNMP and the optional User Datagram Protocol (UDP) port on the remote device. The default is 162.

	Command	Purpose
Step 3	**snmp-server group** *groupname* {**v1** | **v2c** | **v3** {**auth** | **noauth** | **priv**}} [**read** *readview*] [**write** *writeview*] [**notify** *notifyview*] [**access** *access-list*]	Configures a new SNMP group on the remote device. • *groupname*—Specifies the name of the group. • Specify a security model: – **v1** is the least secure of the possible security models. – **v2c** is the second least secure model. It allows transmission of informs and integers twice the normal width. – **v3,** the most secure, requires you to select an authentication level: **auth**—Enables the Message Digest 5 (MD5) and the Secure Hash Algorithm (SHA) packet authentication. **noauth**—Enables the noAuthNoPriv security level. This is the default if no keyword is specified. **priv**—Enables Data Encryption Standard (DES) packet encryption (also called *privacy*). **Note** The **priv** keyword is available only when the cryptographic software image is installed. • (Optional) **read** *readview*—Specifies a string (not to exceed 64 characters) that is the name of the view in which you can only view the contents of the agent. • (Optional) **write** *writeview*—Specifies a string (not to exceed 64 characters) that is the name of the view in which you enter data and configure the contents of the agent. • (Optional) **notify** *notifyview*—Specifies a string (not to exceed 64 characters) that is the name of the view in which you specify a notify, inform, or trap. • (Optional) **access** *access-list*—Specifies a string (not to exceed 64 characters) that is the name of the access list.

	Command	Purpose
Step 4	**snmp-server user** *username groupname* {**remote** *host* [**udp-port** *port*]} {**v1** [**access** *access-list*] | **v2c** [**access** *access-list*] | **v3** [**encrypted**] [**access** *access-list*] [**auth** {**md5** | **sha**} *auth-password*]} [*priv* {**des** | **3des** | **aes** {**128** | **192** | **256**}} *priv-password*]	Adds a new user for an SNMP group. • *username*—Specifies a name of the user on the host that connects to the agent. • *groupname*—Specifies a name of the group to which the user is associated. • **remote**—Specifies a remote SNMP entity to which the user belongs and the hostname or IP address of that entity with the optional UDP port number. The default is 162. • Enters the SNMP version number (**v1**, **v2c**, or **v3**). If you enter **v3**, you have these additional options: – **encrypted**—Specifies that the password appears in encrypted format. This keyword is available only when the **v3** keyword is specified. – **auth**—Specifies an authentication level setting session that can be either the HMAC-MD5-96 (**md5**) or the HMAC-SHA-96 (**sha**) authentication level and requires a password string *auth-password* (not to exceed 64 characters). • If you enter **v3** and the switch is running the cryptographic software image, you can also configure a private (**priv**) encryption algorithm and password string *priv-password* (not to exceed 64 characters). – **priv**—Specifies the User-based Security Model (USM). – **des**—Specifies the use of the 56-bit DES algorithm. – **3des**—Specifies the use of the 168-bit DES algorithm. – **aes**—Specifies the use of the DES algorithm. You must select either 128-bit, 192-bit, or 256-bit encryption. • (Optional) Enters **access** *access-list* with a string (not to exceed 64 characters) that is the name of the access list.
Step 5	**end**	Returns to privileged EXEC mode.

Configuring SNMP Notifications

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**snmp-server engineID remote** *ip-address engineid-string*	Specifies the engine ID for the remote host.

	Command	Purpose
Step 3	**snmp-server user** *username groupname* {**remote** *host* [**udp-port** *port*]} {**v1** [**access** *access-list*] | **v2c** [**access** *access-list*] | **v3** [**encrypted**] [**access** *access-list*] [**auth** {**md5** | **sha**} *auth-password*]}	Configures an SNMP user to be associated with the remote host created in Step 2. **Note** You cannot configure a remote user for an address without first configuring the engine ID for the remote host. Otherwise, you receive an error message, and the command is not executed.
Step 4	**snmp-server group** *groupname* {**v1** | **v2c** | **v3** {**auth** | **noauth** | **priv**}} [**read** *readview*] [**write** *writeview*] [**notify** *notifyview*] [**access** *access-list*]	Configures an SNMP group.
Step 5	**snmp-server host** *host-addr* [**informs** | **traps**] [**version** {**1** | **2c** | **3** {**auth** | **noauth** | **priv**}}] *community-string* [*notification-type*]	Specifies the recipient of an SNMP trap operation. • *host-addr*—Specifies the name or Internet address of the host (the targeted recipient). • (Optional) **informs**—Specifies SNMP informs to be sent to the host. • (Optional) **traps** (the default)—Specifies SNMP traps to be sent to the host. • (Optional) Specifies the SNMP **version** (**1**, **2c**, or **3**). SNMPv1 does not support informs. • (Optional) Version 3—Selects authentication level **auth, noauth**, or **priv**. **Note** The **priv** keyword is available only when the cryptographic software image is installed. • *community-string*—When **version 1** or **version 2c** is specified, enters the password-like community string sent with the notification operation. When **version 3** is specified, enter the SNMPv3 username. **Note** The @ symbol is used for delimiting the context information. Avoid using the @ symbol as part of the SNMP community string when configuring this command. • (Optional) *notification-type*—Specifies a notification type. Use the keywords listed in Table 36-4 on page 36-7. If no type is specified, all notifications are sent.
Step 6	**snmp-server enable traps** *notification-types*	Enables the switch to send traps or informs and specifies the type of notifications to be sent. For a list of notification types, see Table 36-4 on page 36-7, or enter **snmp-server enable traps ?** To enable multiple types of traps, you must enter a separate **snmp-server enable traps** command for each trap type. **Note** When you configure a trap by using the notification type **port-security**, configure the port security trap first, and then configure the port security trap rate: • **snmp-server enable traps port-security** • **snmp-server enable traps port-security trap-rate** *rate*

	Command	Purpose
Step 7	**snmp-server trap-source** *interface-id*	(Optional) Specifies the source interface, which provides the IP address for the trap message. This command also sets the source IP address for informs.
Step 8	**snmp-server queue-length** *length*	(Optional) Establishes the message queue length for each trap host. The range is 1 to 1000; the default is 10.
Step 9	**snmp-server trap-timeout** *seconds*	(Optional) Defines how often to resend trap messages. The range is 1 to 1000; the default is 30 seconds.
Step 10	**end**	Returns to privileged EXEC mode.

Setting the CPU Threshold Notification Types and Values

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**process cpu threshold type** {**total** | **process** | **interrupt**} **rising** *percentage* **interval** *seconds* [**falling** *fall-percentage* **interval** *seconds*]	Sets the CPU threshold notification types and values: • **total**—Sets the notification type to total CPU utilization. • **process**—Sets the notification type to CPU process utilization. • **interrupt**—Sets the notification type to CPU interrupt utilization. • **rising** *percentage*—Specifies the percentage (1 to 100) of CPU resources that, when exceeded for the configured interval, sends a CPU threshold notification. • **interval** *seconds*—Specifies the duration of the CPU threshold violation in seconds (5 to 86400) that, when met, sends a CPU threshold notification. • **falling** *fall-percentage*—Specifies the percentage (1 to 100) of CPU resources that, when usage falls below this level for the configured interval, sends a CPU threshold notification. This value must be equal to or less than the **rising** *percentage* value. If not specified, the **falling** *fall-percentage* value is the same as the **rising** *percentage* value.
Step 3	**end**	Returns to privileged EXEC mode.

Setting the Agent Contact and Location Information

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**snmp-server contact** *text*	Sets the system contact string.
Step 3	**snmp-server location** *text*	Sets the system location string.
Step 4	**end**	Returns to privileged EXEC mode.

Limiting TFTP Servers Used Through SNMP

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**snmp-server tftp-server-list** *access-list-number*	Limits TFTP servers used for configuration file copies through SNMP to the servers in the access list. *access-list-number*—Enters an IP standard access list numbered from 1 to 99 and 1300 to 1999.
Step 3	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*]	Creates a standard access list, repeating the command as many times as necessary. • *access-list-number*—Enters the access list number specified in Step 2. • **deny**—Denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • *source*—Enters the IP address of the TFTP servers that can access the switch. • (Optional) *source-wildcard*—Enters the wildcard bits, in dotted decimal notation, to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 4	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining SNMP

Command	Purpose
show snmp	Displays SNMP statistics.
show snmp engineID [**local** | **remote**]	Displays information on the local SNMP engine and all remote engines that have been configured on the device.
show snmp group	Displays information on each SNMP group on the network.
show snmp pending	Displays information on pending SNMP requests.
show snmp sessions	Displays information on the current SNMP sessions.
show snmp user	Displays information on each SNMP user name in the SNMP users table. **Note** You must use this command to display SNMPv3 configuration information for **auth** | **noauth** | **priv** mode. This information is not displayed in the **show running-config** output.

Configuration Examples for SNMP

Enabling SNMP Versions: Example

This example shows how to enable all versions of SNMP. The configuration permits any SNMP manager to access all objects with read-only permissions using the community string *public*. This configuration does not cause the switch to send any traps.

```
Switch(config)# snmp-server community public
```

Permit SNMP Manager Access: Example

This example shows how to permit any SNMP manager to access all objects with read-only permission using the community string *public*. The switch also sends VTP traps to the hosts 192.180.1.111 and 192.180.1.33 using SNMPv1 and to the host 192.180.1.27 using SNMPv2C. The community string *public* is sent with the traps.

```
Switch(config)# snmp-server community public
Switch(config)# snmp-server enable traps vtp
Switch(config)# snmp-server host 192.180.1.27 version 2c public
Switch(config)# snmp-server host 192.180.1.111 version 1 public
Switch(config)# snmp-server host 192.180.1.33 public
```

Allow Read-Only Access: Example

This example shows how to allow read-only access for all objects to members of access list 4 that use the *comaccess* community string. No other SNMP managers have access to any objects. SNMP Authentication Failure traps are sent by SNMPv2C to the host *cisco.com* using the community string *public*.

```
Switch(config)# snmp-server community comaccess ro 4
Switch(config)# snmp-server enable traps snmp authentication
Switch(config)# snmp-server host cisco.com version 2c public
```

Configure SNMP Traps: Examples

This example shows how to send entity MIB traps to the host *cisco.com*. The community string is restricted. The first line enables the switch to send entity MIB traps in addition to any traps previously enabled. The second line specifies the destination of these traps and overwrites any previous **snmp-server host** commands for the host *cisco.com*.

```
Switch(config)# snmp-server enable traps entity
Switch(config)# snmp-server host cisco.com restricted entity
```

This example shows how to enable the switch to send all traps to the host *myhost.cisco.com* using the community string *public*:

```
Switch(config)# snmp-server enable traps
Switch(config)# snmp-server host myhost.cisco.com public
```

Associating a User with a Remote Host: Example

This example shows how to associate a user with a remote host and to send **auth** (authNoPriv) authentication-level informs when the user enters global configuration mode:

```
Switch(config)# snmp-server engineID remote 192.180.1.27 00000063000100a1c0b4011b
Switch(config)# snmp-server group authgroup v3 auth
Switch(config)# snmp-server user authuser authgroup remote 192.180.1.27 v3 auth md5
mypassword
Switch(config)# snmp-server user authuser authgroup v3 auth md5 mypassword
Switch(config)# snmp-server host 192.180.1.27 informs version 3 auth authuser config
Switch(config)# snmp-server enable traps
Switch(config)# snmp-server inform retries 0
```

Assigning a String to SNMP: Example

This example shows how to assign the string *comaccess* to SNMP, to allow read-only access, and to specify that IP access list 4 can use the community string to gain access to the switch SNMP agent:

```
Switch(config)# snmp-server community comaccess ro 4
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
Cisco IOS SNMP syntax and usage	*Cisco IOS Network Management Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 37

Configuring Network Security with ACLs

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Restrictions for Network Security with ACLs

The switch does not support these Cisco IOS router ACL-related features:

- Non-IP protocol ACLs (see Table 37-1 on page 37-5) or bridge-group ACLs
- IP accounting
- Inbound and outbound rate limiting (except with QoS ACLs)
- Reflexive ACLs or dynamic ACLs (except for some specialized dynamic ACLs used by the switch clustering feature)
- ACL logging for port ACLs and VLAN maps

Information About Network Security with ACLs

ACLs

Packet filtering can help limit network traffic and restrict network use by certain users or devices. ACLs filter traffic as it passes through a router or switch and permit or deny packets crossing specified interfaces or VLANs. An ACL is a sequential collection of permit and deny conditions that apply to packets. When a packet is received on an interface, the switch compares the fields in the packet against any applied ACLs to verify that the packet has the required permissions to be forwarded, based on the criteria specified in the access lists. One by one, it tests packets against the conditions in an access list. The first match decides whether the switch accepts or rejects the packets. Because the switch stops testing after the first match, the order of conditions in the list is critical. If no conditions match, the

switch rejects the packet. If there are no restrictions, the switch forwards the packet; otherwise, the switch drops the packet. The switch can use ACLs on all packets it forwards, including packets bridged within a VLAN.

You configure access lists on a router or Layer 3 switch to provide basic security for your network. If you do not configure ACLs, all packets passing through the switch could be allowed onto all parts of the network. You can use ACLs to control which hosts can access different parts of a network or to decide which types of traffic are forwarded or blocked at router interfaces. For example, you can allow e-mail traffic to be forwarded but not Telnet traffic. ACLs can be configured to block inbound traffic, outbound traffic, or both.

An ACL contains an ordered list of access control entries (ACEs). Each ACE specifies *permit* or *deny* and a set of conditions the packet must satisfy in order to match the ACE. The meaning of *permit* or *deny* depends on the context in which the ACL is used.

The switch supports IP ACLs and Ethernet (MAC) ACLs:

- IP ACLs filter IPv4 traffic, including TCP, User Datagram Protocol (UDP), Internet Group Management Protocol (IGMP), and Internet Control Message Protocol (ICMP).
- Ethernet ACLs filter non-IP traffic.

This switch also supports quality of service (QoS) classification ACLs. For more information, see the "Classification Based on QoS ACLs" section on page 38-13.

These sections contain this conceptual information:

- Supported ACLs, page 37-2
- Handling Fragmented and Unfragmented Traffic, page 37-3

Supported ACLs

Port ACLs access-control traffic entering a Layer 2 interface. The switch does not support port ACLs in the outbound direction. You can apply only one IP access list and one MAC access list to a Layer 2 interface. For more information, see the "Port ACLs" section on page 37-2.

If IEEE 802.1Q tunneling is configured on an interface, any IEEE 802.1Q encapsulated IP packets received on the tunnel port can be filtered by MAC ACLs, but not by IP ACLs. This is because the switch does not recognize the protocol inside the IEEE 802.1Q header. This restriction applies to router ACLs and port ACLs.

Port ACLs

Note To use this feature, the switch must be running the LAN Base image.

Port ACLs are ACLs that are applied to Layer 2 interfaces on a switch. Port ACLs are supported only on physical interfaces and not on EtherChannel interfaces and can be applied only on interfaces in the inbound direction. These access lists are supported:

- Standard IP access lists using source addresses
- Extended IP access lists using source and destination addresses and optional protocol type information
- MAC extended access lists using source and destination MAC addresses and optional protocol type information

The switch examines ACLs associated with all inbound features configured on a given interface and permits or denies packet forwarding based on how the packet matches the entries in the ACL. In this way, ACLs control access to a network or to part of a network. Figure 37-1 is an example of using port ACLs to control access to a network when all workstations are in the same VLAN. ACLs applied at the Layer 2 input would allow Host A to access the Human Resources network, but prevent Host B from accessing the same network. Port ACLs can only be applied to Layer 2 interfaces in the inbound direction.

Figure 37-1 Using ACLs to Control Traffic to a Network

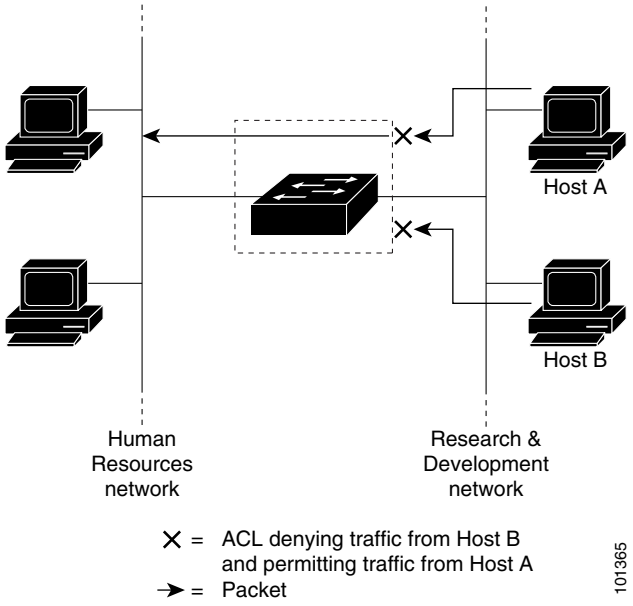

When you apply a port ACL to a trunk port, the ACL filters traffic on all VLANs present on the trunk port. When you apply a port ACL to a port with voice VLAN, the ACL filters traffic on both data and voice VLANs.

With port ACLs, you can filter IP traffic by using IP access lists and non-IP traffic by using MAC addresses. You can filter both IP and non-IP traffic on the same Layer 2 interface by applying both an IP access list and a MAC access list to the interface.

Note You cannot apply more than one IP access list and one MAC access list to a Layer 2 interface. If an IP access list or MAC access list is already configured on a Layer 2 interface and you apply a new IP access list or MAC access list to the interface, the new ACL replaces the previously configured one.

Handling Fragmented and Unfragmented Traffic

IP packets can be fragmented as they cross the network. When this happens, only the fragment containing the beginning of the packet contains the Layer 4 information, such as TCP or UDP port numbers, ICMP type and code, and so on. All other fragments are missing this information.

Some ACEs do not check Layer 4 information and therefore can be applied to all packet fragments. ACEs that do test Layer 4 information cannot be applied in the standard manner to most of the fragments in a fragmented IP packet. When the fragment contains no Layer 4 information and the ACE tests some Layer 4 information, the matching rules are modified:

- Permit ACEs that check the Layer 3 information in the fragment (including protocol type, such as TCP, UDP, and so on) are considered to match the fragment regardless of what the missing Layer 4 information might have been.
- Deny ACEs that check Layer 4 information never match a fragment unless the fragment contains Layer 4 information.

Consider access list 102, configured with these commands, applied to three fragmented packets:

```
Switch(config)# access-list 102 permit tcp any host 10.1.1.1 eq smtp
Switch(config)# access-list 102 deny tcp any host 10.1.1.2 eq telnet
Switch(config)# access-list 102 permit tcp any host 10.1.1.2
Switch(config)# access-list 102 deny tcp any any
```

Note In the first and second ACEs in the examples, the *eq* keyword after the destination address means to test for the TCP-destination-port well-known numbers equaling Simple Mail Transfer Protocol (SMTP) and Telnet, respectively.

- Packet A is a TCP packet from host 10.2.2.2., port 65000, going to host 10.1.1.1 on the SMTP port. If this packet is fragmented, the first fragment matches the first ACE (a permit) as if it were a complete packet because all Layer 4 information is present. The remaining fragments also match the first ACE, even though they do not contain the SMTP port information, because the first ACE only checks Layer 3 information when applied to fragments. The information in this example is that the packet is TCP and that the destination is 10.1.1.1.
- Packet B is from host 10.2.2.2, port 65001, going to host 10.1.1.2 on the Telnet port. If this packet is fragmented, the first fragment matches the second ACE (a deny) because all Layer 3 and Layer 4 information is present. The remaining fragments in the packet do not match the second ACE because they are missing Layer 4 information. Instead, they match the third ACE (a permit).

  Because the first fragment was denied, host 10.1.1.2 cannot reassemble a complete packet, so packet B is effectively denied. However, the later fragments that are permitted will consume bandwidth on the network and resources of host 10.1.1.2 as it tries to reassemble the packet.
- Fragmented packet C is from host 10.2.2.2, port 65001, going to host 10.1.1.3, port ftp. If this packet is fragmented, the first fragment matches the fourth ACE (a deny). All other fragments also match the fourth ACE because that ACE does not check any Layer 4 information and because Layer 3 information in all fragments shows that they are being sent to host 10.1.1.3, and the earlier permit ACEs were checking different hosts.

IPv4 ACLs

Configuring IPv4 ACLs on the switch is the same as configuring IPv4 ACLs on other Cisco switches and routers.

Step 1 Create an ACL by specifying an access list number or name and the access conditions.

Step 2 Apply the ACL to interfaces or terminal lines.

Standard and Extended IPv4 ACLs

This section describes IP ACLs. An ACL is a sequential collection of permit and deny conditions. One by one, the switch tests packets against the conditions in an access list. The first match determines whether the switch accepts or rejects the packet. Because the switch stops testing after the first match, the order of the conditions is critical. If no conditions match, the switch denies the packet.

The software supports these types of ACLs or access lists for IPv4:

- Standard IP access lists use source addresses for matching operations.
- Extended IP access lists use source and destination addresses for matching operations and optional protocol-type information for finer granularity of control.

The switch always rewrites the order of standard access lists so that entries with **host** matches and entries with matches having a *don't care* mask of 0.0.0.0 are moved to the top of the list, above any entries with non-zero *don't care* masks. Therefore, in **show** command output and in the configuration file, the ACEs do not necessarily appear in the order in which they were entered.

After creating a numbered standard IPv4 ACL, you can apply it to terminal lines (see the "Applying an IPv4 ACL to a Terminal Line" section on page 37-17), to interfaces (see the "Applying an IPv4 ACL to an Interface" section on page 37-17), or to VLANs (see the "Monitoring and Maintaining Network Security with ACLs" section on page 37-19).

Access List Numbers

The number you use to denote your ACL shows the type of access list that you are creating. Table 37-1 lists the access-list number and corresponding access list type and shows whether or not they are supported in the switch. The switch supports IPv4 standard and extended access lists, numbers 1 to 199 and 1300 to 2699.

Table 37-1 Access List Numbers

Access List Number	Type	Supported
1–99	IP standard access list	Yes
100–199	IP extended access list	Yes
200–299	Protocol type-code access list	No
300–399	DECnet access list	No
400–499	XNS standard access list	No
500–599	XNS extended access list	No
600–699	AppleTalk access list	No
700–799	48-bit MAC address access list	No
800–899	IPX standard access list	No
900–999	IPX extended access list	No
1000–1099	IPX SAP access list	No
1100–1199	Extended 48-bit MAC address access list	No
1200–1299	IPX summary address access list	No
1300–1999	IP standard access list (expanded range)	Yes
2000–2699	IP extended access list (expanded range)	Yes

Note In addition to numbered standard and extended ACLs, you can also create standard and extended named IP ACLs by using the supported numbers. That is, the name of a standard IP ACL can be 1 to 99; the name of an extended IP ACL can be 100 to 199. The advantage of using named ACLs instead of numbered lists is that you can delete individual entries from a named list.

ACL Logging

The switch software can provide logging messages about packets permitted or denied by a standard IP access list. That is, any packet that matches the ACL causes an informational logging message about the packet to be sent to the console. The level of messages logged to the console is controlled by the logging console commands controlling the syslog messages.

Note Because routing is done in hardware and logging is done in software, if a large number of packets match a *permit* or *deny* ACE containing a **log** keyword, the software might not be able to match the hardware processing rate, and not all packets will be logged.

The first packet that triggers the ACL causes a logging message right away, and subsequent packets are collected over 5-minute intervals before they appear or logged. The logging message includes the access list number, whether the packet was permitted or denied, the source IP address of the packet, and the number of packets from that source permitted or denied in the prior 5-minute interval.

Numbered Extended ACL

Although standard ACLs use only source addresses for matching, you can use extended ACL source and destination addresses for matching operations and optional protocol type information for finer granularity of control. When you are creating ACEs in numbered extended access lists, remember that after you create the ACL, any additions are placed at the end of the list. You cannot reorder the list or selectively add or remove ACEs from a numbered list.

Some protocols also have specific parameters and keywords that apply to that protocol.

These IP protocols are supported (protocol keywords are in parentheses in bold):

- Authentication Header Protocol (**ahp**)
- Enhanced Interior Gateway Routing Protocol (**eigrp)**
- Encapsulation Security Payload (**esp**)
- generic routing encapsulation (**gre**)
- Internet Control Message Protocol (**icmp)**
- Internet Group Management Protocol (**igmp**)
- any Interior Protocol (**ip**)
- IP in IP tunneling (**ipinip**)
- KA9Q NOS-compatible IP over IP tunneling (**nos**)
- Open Shortest Path First routing (**ospf**)
- Payload Compression Protocol (**pcp**)
- Protocol Independent Multicast (**pim**)
- Transmission Control Protocol (**tcp**)

- User Datagram Protocol (**udp**)

Note ICMP echo-reply cannot be filtered. All other ICMP codes or types can be filtered.

Note The switch does not support dynamic or reflexive access lists. It also does not support filtering based on the type of service (ToS) minimize-monetary-cost bit.

Supported parameters can be grouped into these categories: TCP, UDP, ICMP, IGMP, or other IP.

After an ACL is created, any additions (possibly entered from the terminal) are placed at the end of the list. You cannot selectively add or remove access list entries from a numbered access list.

Note When you are creating an ACL, remember that, by default, the end of the access list contains an implicit deny statement for all packets if it did not find a match before reaching the end.

After creating a numbered extended ACL, you can apply it to terminal lines (see the "Applying an IPv4 ACL to a Terminal Line" section on page 37-17), to interfaces (see the "Applying an IPv4 ACL to an Interface" section on page 37-17), or to VLANs (see the "Monitoring and Maintaining Network Security with ACLs" section on page 37-19).

Resequencing ACEs in an ACL

Sequence numbers for the entries in an access list are automatically generated when you create a new ACL. You can use the **ip access-list resequence** global configuration command to edit the sequence numbers in an ACL and change the order in which ACEs are applied. For example, if you add a new ACE to an ACL, it is placed at the bottom of the list. By changing the sequence number, you can move the ACE to a different position in the ACL.

Named Standard and Extended ACLs

You can identify IPv4 ACLs with an alphanumeric string (a name) rather than a number. You can use named ACLs to configure more IPv4 access lists in a router than if you were to use numbered access lists. If you identify your access list with a name rather than a number, the mode and command syntax are slightly different. However, not all commands that use IP access lists accept a named access list.

Note The name you give to a standard or extended ACL can also be a number in the supported range of access list numbers. That is, the name of a standard IP ACL can be 1 to 99; the name of an extended IP ACL can be 100 to 199. The advantage of using named ACLs instead of numbered lists is that you can delete individual entries from a named list.

Consider these guidelines and limitations before configuring named ACLs:

- Not all commands that accept a numbered ACL accept a named ACL. ACLs for packet filters and route filters on interfaces can use a name.
- A standard ACL and an extended ACL cannot have the same name.
- Numbered ACLs are also available, as described in the "Creating a Numbered Standard ACL" section on page 37-11.

When you are creating standard extended ACLs, remember that, by default, the end of the ACL contains an implicit deny statement for everything if it did not find a match before reaching the end. For standard ACLs, if you omit the mask from an associated IP host address access list specification, 0.0.0.0 is assumed to be the mask.

After you create an ACL, any additions are placed at the end of the list. You cannot selectively add ACL entries to a specific ACL. However, you can use **no permit** and **no deny** access-list configuration mode commands to remove entries from a named ACL. This example shows how you can delete individual ACEs from the named access list *border-list*:

```
Switch(config)# ip access-list extended border-list
Switch(config-ext-nacl)# no permit ip host 10.1.1.3 any
```

Being able to selectively remove lines from a named ACL is one reason you might use named ACLs instead of numbered ACLs.

Time Ranges with ACLs

You can selectively apply extended ACLs based on the time of day and the week by using the **time-range** global configuration command. First, define a time-range name and set the times and the dates or the days of the week in the time range. Then enter the time-range name when applying an ACL to set restrictions to the access list. You can use the time range to define when the permit or deny statements in the ACL are in effect, for example, during a specified time period or on specified days of the week.

These are some of the many possible benefits of using time ranges:

- You have more control over permitting or denying a user access to resources, such as an application (identified by an IP address/mask pair and a port number).
- You can control logging messages. ACL entries can be set to log traffic only at certain times of the day. Therefore, you can simply deny access without needing to analyze many logs generated during peak hours.

Time-based access lists trigger CPU activity because the new configuration of the access list must be merged with other features and the combined configuration loaded into the TCAM. For this reason, you should be careful not to have several access lists configured to take affect in close succession (within a small number of minutes of each other.)

Note The time range relies on the switch system clock; therefore, you need a reliable clock source. We recommend that you use Network Time Protocol (NTP) to synchronize the switch clock. For more information, see the "System Time and Date Management" section on page 7-1.

Comments in ACLs

You can use the **remark** keyword to include comments (remarks) about entries in any IP standard or extended ACL. The remarks make the ACL easier for you to understand and scan. Each remark line is limited to 100 characters.

The remark can go before or after a permit or deny statement. You should be consistent about where you put the remark so that it is clear which remark describes which permit or deny statement. For example, it would be confusing to have some remarks before the associated permit or deny statements and some remarks after the associated statements.

To include a comment for IP numbered standard or extended ACLs, use the **access-list** *access-list number* **remark** *remark* global configuration command. To remove the remark, use the **no** form of this command.

IPv4 ACL to a Terminal Line

You can use numbered ACLs to control access to one or more terminal lines. You cannot apply named ACLs to lines. You must set identical restrictions on all the virtual terminal lines because a user can attempt to connect to any of them.

For procedures for applying ACLs to interfaces, see the "Applying an IPv4 ACL to an Interface" section on page 37-17. For applying ACLs to VLANs, see the "Monitoring and Maintaining Network Security with ACLs" section on page 37-19.

IPv4 ACL Application to an Interface Guidelines

- Apply an ACL only to inbound Layer 2 ports.
- Apply an ACL to either outbound or inbound Layer 3 interfaces.
- When controlling access to an interface, you can use a named or numbered ACL.
- If you apply an ACL to a port that is a member of a VLAN, the port ACL takes precedence over an ACL applied to the VLAN interface.
- If you apply an ACL to a Layer 2 interface that is a member of a VLAN, the Layer 2 (port) ACL takes precedence over an input Layer 3 ACL applied to the VLAN interface. The port ACL always filters incoming packets received on the Layer 2 port.
- If you apply an ACL to a Layer 3 interface and routing is not enabled, the ACL only filters packets that are intended for the CPU, such as SNMP, Telnet, or web traffic. You do not have to enable routing to apply ACLs to Layer 2 interfaces.
- When private VLANs are configured, you can apply router ACLs only on the primary-VLAN SVIs. The ACL is applied to both primary and secondary VLAN Layer 3 traffic.

Note By default, the router sends Internet Control Message Protocol (ICMP) unreachable messages when a packet is denied by an access group. These access-group denied packets are not dropped in hardware but are bridged to the switch CPU so that it can generate the ICMP-unreachable message. Port ACLs are an exception. They do not generate ICMP unreachable messages.

ICMP unreachable messages can be disabled on router ACLs with the **no ip unreachables** interface command.

For inbound ACLs, after receiving a packet, the switch checks the packet against the ACL. If the ACL permits the packet, the switch continues to process the packet. If the ACL rejects the packet, the switch discards the packet.

For outbound ACLs, after receiving and sending a packet to a controlled interface, the switch checks the packet against the ACL. If the ACL permits the packet, the switch sends the packet. If the ACL rejects the packet, the switch discards the packet.

By default, the input interface sends ICMP Unreachable messages whenever a packet is discarded, regardless of whether the packet was discarded because of an ACL on the input interface or because of an ACL on the output interface. ICMP Unreachables are normally limited to no more than one every one-half second per input interface, but this can be changed by using the **ip icmp rate-limit unreachable** global configuration command.

When you apply an undefined ACL to an interface, the switch acts as if the ACL has not been applied to the interface and permits all packets. Remember this behavior if you use undefined ACLs for network security.

Hardware and Software Handling of IP ACLs

ACL processing is primarily accomplished in hardware, but requires forwarding of some traffic flows to the CPU for software processing. If the hardware reaches its capacity to store ACL configurations, packets are sent to the CPU for forwarding. The forwarding rate for software-forwarded traffic is substantially less than for hardware-forwarded traffic.

Note If an ACL configuration cannot be implemented in hardware due to an out-of-resource condition on a switch, then only the traffic in that VLAN arriving on that switch is affected (forwarded in software). Software forwarding of packets might adversely impact the performance of the switch, depending on the number of CPU cycles that this consumes.

For router ACLs, other factors can cause packets to be sent to the CPU:

- Using the **log** keyword
- Generating ICMP unreachable messages

When traffic flows are both logged and forwarded, forwarding is done by hardware, but logging must be done by software. Because of the difference in packet handling capacity between hardware and software, if the sum of all flows being logged (both permitted flows and denied flows) is of significant bandwidth, not all of the packets that are forwarded can be logged.

If router ACL configuration cannot be applied in hardware, packets arriving in a VLAN that must be routed are routed in software, but are bridged in hardware. If ACLs cause large numbers of packets to be sent to the CPU, the switch performance can be negatively affected.

When you enter the **show ip access-lists** privileged EXEC command, the match count displayed does not account for packets that are access controlled in hardware. Use the **show access-lists hardware counters** privileged EXEC command to obtain some basic hardware ACL statistics for switched and routed packets.

Troubleshooting ACLs

If this ACL manager message appears, where [chars] is the access-list name, the switch then has insufficient resources to create a hardware representation of the ACL.

```
ACLMGR-2-NOVMR: Cannot generate hardware representation of access list [chars]
```

The resources include hardware memory and label space but not CPU memory. A lack of available logical operation units or specialized hardware resources causes this problem. Logical operation units are needed for a TCP flag match or a test other than **eq** (**ne**, **gt**, **lt**, or **range**) on TCP, UDP, or SCTP port numbers.

Use one of these workarounds:

- Modify the ACL configuration to use fewer resources.
- Rename the ACL with a name or number that alphanumerically precedes the ACL names or numbers.

To determine the specialized hardware resources, enter the **show platform layer4 acl map** privileged EXEC command. If the switch does not have available resources, the output shows that index 0 to index 15 are not available.

For more information about configuring ACLs with insufficient resources, see CSCsq63926 in the Bug Toolkit.

Named MAC Extended ACLs

You can filter non-IPv4 traffic on a VLAN or on a Layer 2 interface by using MAC addresses and named MAC extended ACLs. The procedure is similar to that of configuring other extended named ACLs.

Note You cannot apply named MAC extended ACLs to Layer 3 interfaces.

Note Though visible in the command-line help strings, **appletalk** is not supported as a matching condition for the **deny** and **permit** MAC access-list configuration mode commands.

MAC ACL to a Layer 2 Interface

After you create a MAC ACL, you can apply it to a Layer 2 interface to filter non-IP traffic coming in that interface. When you apply the MAC ACL, consider these guidelines:

- If you apply an ACL to a Layer 2 interface that is a member of a VLAN, the Layer 2 (port) ACL takes precedence over an input Layer 3 ACL applied to the VLAN interface. Incoming packets received on the Layer 2 port are always filtered by the port ACL.
- You can apply no more than one IP access list and one MAC access list to the same Layer 2 interface. The IP access list filters only IP packets, and the MAC access list filters non-IP packets.
- A Layer 2 interface can have only one MAC access list. If you apply a MAC access list to a Layer 2 interface that has a MAC ACL configured, the new ACL replaces the previously configured one.

How to Configure Network Security with ACLs

Creating a Numbered Standard ACL

Note When creating an ACL, remember that, by default, the end of the ACL contains an implicit deny statement for all packets that it did not find a match for before reaching the end. With standard access lists, if you omit the mask from an associated IP host address ACL specification, 0.0.0.0 is assumed to be the mask.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*] [**log**]	Defines a standard IPv4 access list by using a source address and wildcard. *access-list-number*—Specifies a decimal number from 1 to 99 or 1300 to 1999. **deny** or **permit**—Specifies whether to deny or permit access if conditions are matched. *source*—Specifies the source address of the network or host from which the packet is being sent specified as: • The 32-bit quantity in dotted-decimal format. • The keyword **any** as an abbreviation for *source* and *source-wildcard* of 0.0.0.0 255.255.255.255. You do not need to enter a source-wildcard. • The keyword **host** as an abbreviation for source and source-wildcard of *source* 0.0.0.0. (Optional) *source-wildcard*—Applies wildcard bits to the source. (Optional) **log**—Causes an informational logging message about the packet that matches the entry to be sent to the console.
Step 3	**end**	Returns to privileged EXEC mode.

Creating a Numbered Extended ACL

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2a	**access-list** *access-list-number* {**deny** | **permit**} *protocol source source-wildcard destination destination-wildcard* [**precedence** *precedence*] [**tos** *tos*] [**fragments**] [**log**] [**log-input**] [**time-range** *time-range-name*] [**dscp** *dscp*] **Note** If you enter a **dscp** value, you cannot enter **tos** or **precedence**. You can enter both a **tos** and a **precedence** value with no **dscp**.	Defines an extended IPv4 access list and the access conditions. *access-list-number*—Specifies a decimal number from 100 to 199 or 2000 to 2699. **deny** or **permit**—Specifies whether to deny or permit the packet if conditions are matched. *protocol*—Specifies the name or number of an IP protocol: **ahp**, **eigrp**, **esp**, **gre**, **icmp**, **igmp**, **igrp**, **ip**, **ipinip**, **nos**, **ospf**, **pcp**, **pim**, **tcp**, or **udp**, or an integer in the range 0 to 255 representing an IP protocol number. To match any Internet protocol (including ICMP, TCP, and UDP), use the keyword **ip**. **Note** This step includes options for most IP protocols. For additional specific parameters for TCP, UDP, ICMP, and IGMP, see steps 2b through 2e. *source*—The number of the network or host from which the packet is sent. *source-wildcard*—Applies wildcard bits to the source. *destination*—The network or host number to which the packet is sent. *destination-wildcard*—Applies wildcard bits to the destination. *source*, *source-wildcard*, *destination*, and *destination-wildcard* can be specified as: • The 32-bit quantity in dotted-decimal format. • The keyword **any** for 0.0.0.0 255.255.255.255 (any host). • The keyword **host** for a single host 0.0.0.0. The other keywords are optional and have these meanings: • **precedence**—Matches packets with a precedence level specified as a number from 0 to 7 or by name: **routine** (**0**), ***priority*** (**1**), **immediate** (**2**), **flash** (**3**), **flash-override** (**4**), **critical** (**5**), **internet** (6), **network** (**7**). • **fragments**—Checks noninitial fragments. • **tos**—Matches by type of service level, specified by a number from 0 to 15 or a name: **normal** (**0**), **max-reliability** (**2**), **max-throughput** (**4**), **min-delay** (**8**). • **log**—Creates an informational logging message to be sent to the console about the packet that matches the entry or **log-input** to include the input interface in the log entry. • **time-range**—For an explanation of this keyword, see the "Using Time Ranges with ACLs" section on page 37-16. • **dscp**—Matches packets with the DSCP value specified by a number from 0 to 63, or use the question mark (?) to see a list of available values.

	Command	Purpose
or	**access-list** *access-list-number* {**deny** | **permit**} *protocol* **any any** [**precedence** *precedence*] [**tos** *tos*] [**fragments**] [**log**] [**log-input**] [**time-range** *time-range-name*] [**dscp** *dscp*]	In access-list configuration mode, defines an extended IP access list using an abbreviation for a source and source wildcard of 0.0.0.0 255.255.255.255 and an abbreviation for a destination and destination wildcard of 0.0.0.0 255.255.255.255. You can use the **any** keyword in place of source and destination address and wildcard.
or	**access-list** *access-list-number* {**deny** | **permit**} *protocol* **host** *source* **host** *destination* [**precedence** *precedence*] [**tos** *tos*] [**fragments**] [**log**] [**log-input**] [**time-range** *time-range-name*] [**dscp** *dscp*]	Defines an extended IP access list by using an abbreviation for a source and a source wildcard of *source* 0.0.0.0 and an abbreviation for a destination and destination wildcard of *destination* 0.0.0.0. You can use the **host** keyword in place of the source and destination wildcard or mask.
Step 2b	**access-list** *access-list-number* {**deny** | **permit**} **tcp** *source source-wildcard* [*operator port*] *destination destination-wildcard* [*operator port*] [**established**] [**precedence** *precedence*] [**tos** *tos*] [**fragments**] [**log**] [**log-input**] [**time-range** *time-range-name*] [**dscp** *dscp*] [*flag*]	(Optional) Defines an extended TCP access list and the access conditions. Enter **tcp** for Transmission Control Protocol. The parameters are the same as those described in Step 2a, with these exceptions: (Optional) *operator* and *port* compare source (if positioned after *source source-wildcard*) or destination (if positioned after *destination destination-wildcard*) port. Possible operators include **eq** (equal), **gt** (greater than), **lt** (less than), **neq** (not equal), and **range** (inclusive range). Operators require a port number (**range** requires two port numbers separated by a space). *port* number is a decimal number (from 0 to 65535) or the name of a TCP port. To see TCP port names, use the ? or see the "Configuring IP Services" section in the "IP Addressing and Services" chapter of the *Cisco IOS IP Configuration Guide, Release 12.2*. Use only TCP port numbers or names when filtering TCP. The other optional keywords have these meanings: • **established**—Matches an established connection. This has the same function as matching on the **ack** or **rst** flag. • *flag*—Matches one of these flags by the specified TCP header bits: **ack** (acknowledge), **fin** (finish), **psh** (push), **rst** (reset), **syn** (synchronize), or **urg** (urgent).
Step 2c	**access-list** *access-list-number* {**deny** | **permit**} **udp** *source source-wildcard* [*operator port*] *destination destination-wildcard* [*operator port*] [**precedence** *precedence*] [**tos** *tos*] [**fragments**] [**log**] [**log-input**] [**time-range** *time-range-name*] [**dscp** *dscp*]	(Optional) Defines an extended UDP access list and the access conditions. **udp**—The User Datagram Protocol. The UDP parameters are the same as those described for TCP except that the [*operator* [*port*]] port number or name must be a UDP port number or name, and the **flag** and **established** parameters are not valid for UDP.

	Command	Purpose
Step 2d	**access-list** *access-list-number* {**deny** | **permit**} **icmp** *source source-wildcard destination destination-wildcard* [*icmp-type* | [[*icmp-type icmp-code*] | [*icmp-message*]] [**precedence** *precedence*] [**tos** *tos*] [**fragments**] [**log**] [**log-input**] [**time-range** *time-range-name*] [**dscp** *dscp*]	(Optional) Defines an extended ICMP access list and the access conditions. **icmp**—Internet Control Message Protocol. The ICMP parameters are the same as those described for most IP protocols in Step 2a, with the addition of the ICMP message type and code parameters. These optional keywords have these meanings: • *icmp-type*—Filters by ICMP message type, a number from 0 to 255. • *icmp-code*—Filters ICMP packets that are filtered by the ICMP message code type, a number from 0 to 255. • *icmp-message*—Filters ICMP packets by the ICMP message type name or the ICMP message type and code name. To see a list of ICMP message type names and code names, use the ?, or see the "Configuring IP Services" section of the *Cisco IOS IP Configuration Guide, Release 12.2*.
Step 2e	**access-list** *access-list-number* {**deny** | **permit**} **igmp** *source source-wildcard destination destination-wildcard* [*igmp-type*] [**precedence** *precedence*] [**tos** *tos*] [**fragments**] [**log**] [**log-input**] [**time-range** *time-range-name*] [**dscp** *dscp*]	(Optional) Defines an extended IGMP access list and the access conditions. **igmp**—Internet Group Management Protocol. The IGMP parameters are the same as those described for most IP protocols in Step 2a, with this optional parameter. *igmp-type*—Matches IGMP message type, enters a number from 0 to 15, or enters the message name (**dvmrp**, **host-query**, **host-report**, **pim,** or **trace**).
Step 3	**end**	Returns to privileged EXEC mode.

Creating Named Standard and Extended ACLs

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip access-list standard** *name* or **ip access-list extended** *name*	Defines a standard IPv4 access list using a name, and enters access-list configuration mode. The name can be a number from 1 to 99. or Defines an extended IPv4 access list using a name, and enters access-list configuration mode. The name can be a number from 100 to 199.

	Command	Purpose
Step 3	**{deny | permit}** {*source* [*source-wildcard*] | **host** *source* | **any**} [**log**] or {**deny** | **permit**} *protocol* {*source* [*source-wildcard*] | **host** *source* | **any**} {*destination* [*destination-wildcard*] | **host** *destination* | **any**} [**precedence** *precedence*] [**tos** *tos*] [**established**] [**log**] [**time-range** *time-range-name*]	In access-list configuration mode, specifies one or more conditions denied or permitted to decide if the packet is forwarded or dropped. • **host** *source*—A source and source wildcard of *source* 0.0.0.0. • **any**—A source and source wildcard of 0.0.0.0 255.255.255.255. or In access-list configuration mode, specify the conditions allowed or denied. Use the **log** keyword to get access list logging messages, including violations. See the "Creating a Numbered Extended ACL" section on page 37-13 for definitions of protocols and other keywords. • **host** *source*—A source and source wildcard of *source* 0.0.0.0. • ***host*** *destination*—A destination and destination wildcard of *destination* 0.0.0.0. • **any**—A source and source wildcard or destination and destination wildcard of 0.0.0.0 255.255.255.255.
Step 4	**end**	Returns to privileged EXEC mode.

Using Time Ranges with ACLs

Repeat the steps if you have multiple items that you want in effect at different times.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**time-range** *time-range-name*	Assigns a meaningful name (for example, *workhours*) to the time range to be created, and enters time-range configuration mode. The name cannot contain a space or quotation mark and must begin with a letter.
Step 3	**absolute** [**start** *time date*] [**end** *time date*] or **periodic** *day-of-the-week hh:mm to* [*day-of-the-week*] *hh:mm* or **periodic** {**weekdays** | **weekend** | **daily**} *hh:mm to hh:mm*	Specifies when the function it will be applied to is operational. • You can use only one **absolute** statement in the time range. If you configure more than one absolute statement, only the one configured last is executed. • You can enter multiple **periodic** statements. For example, you could configure different hours for weekdays and weekends. See the example configurations.
Step 4	**end**	Returns to privileged EXEC mode.

Applying an IPv4 ACL to a Terminal Line

This task restricts incoming and outgoing connections between a virtual terminal line and the addresses in an ACL:

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**line** [**console** | **vty**] *line-number*	Identifies a specific line to configure, and enters in-line configuration mode. • **console**—Specifies the console terminal line. The console port is DCE. • **vty**—Specifies a virtual terminal for remote console access. The *line-number* is the first line number in a contiguous group that you want to configure when the line type is specified. The range is from 0 to 16.
Step 3	**access-class** *access-list-number* {**in** | **out**}	Restricts incoming and outgoing connections between a particular virtual terminal line (into a device) and the addresses in an access list.
Step 4	**end**	Returns to privileged EXEC mode.

Applying an IPv4 ACL to an Interface

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Identifies a specific interface for configuration, and enters interface configuration mode. The interface is a Layer 2 interface (port ACL).
Step 3	**ip access-group** {*access-list-number* | *name*} {**in** | **out**}	Controls access to the specified interface. The **out** keyword is not supported for Layer 2 interfaces (port ACLs).
Step 4	**end**	Returns to privileged EXEC mode.

Creating Named MAC Extended ACLs

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mac access-list extended** *name*	Defines an extended MAC access list using a name.

	Command	Purpose
Step 3	{**deny** | **permit**} {**any** | **host** *source MAC address* | *source MAC address mask*} {**any** | **host** *destination MAC address* | *destination MAC address mask*} [*type mask* | **lsap** *lsap mask* | **aarp** | **amber** | **dec-spanning** | **decnet-iv** | **diagnostic** | **dsm** | **etype-6000** | **etype-8042** | **lat** | **lavc-sca** | **mop-console** | **mop-dump** | **msdos** | **mumps** | **netbios** | **vines-echo** |**vines-ip** | **xns-idp** | *0-65535*] [**cos** *cos*]	In extended MAC access-list configuration mode, specifies to **permit** or **deny any** source MAC address, a source MAC address with a mask, or a specific **host** source MAC address and **any** destination MAC address, destination MAC address with a mask, or a specific destination MAC address. (Optional) You can also enter these options: • *type mask*—Specifies an arbitrary EtherType number of a packet with Ethernet II or SNAP encapsulation in decimal, hexadecimal, or octal with optional mask of *don't care* bits applied to the EtherType before testing for a match. • **lsap** *lsap mask*—Specifies an LSAP number of a packet with IEEE 802.2 encapsulation in decimal, hexadecimal, or octal with optional mask of *don't care* bits. • **aarp** | **amber** | **dec-spanning** | **decnet-iv** | **diagnostic** | **dsm** | **etype-6000** | **etype-8042** | **lat** | **lavc-sca** | **mop-console** | **mop-dump** | **msdos** | **mumps** | **netbios** | **vines-echo** |**vines-ip** | **xns-idp**—Specifies a non-IP protocol. • **cos** *cos*—Specifies an IEEE 802.1Q cost of service number from 0 to 7 used to set priority.
Step 4	**end**	Returns to privileged EXEC mode.

Applying a MAC ACL to a Layer 2 Interface

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Identifies a specific interface, and enters interface configuration mode. The interface must be a physical Layer 2 interface (port ACL).
Step 3	**mac access-group** {*name*} {**in**}	Controls access to the specified interface by using the MAC access list. Port ACLs are supported only in the inbound direction.
Step 4	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining Network Security with ACLs

Command	Purpose
show access-lists [*number* | *name*]	Displays the contents of one or all current IP and MAC address access lists or a specific access list (numbered or named).
show ip access-lists [*number* | *name*]	Displays the contents of all current IP access lists or a specific IP access list (numbered or named).
show ip interface *interface-id*	Displays detailed configuration and status of an interface. If IP is enabled on the interface and ACLs have been applied by using the **ip access-group** interface configuration command, the access groups are included in the display.
show running-config [**interface** *interface-id*]	Displays the contents of the configuration file for the switch or the specified interface, including all configured MAC and IP access lists and which access groups are applied to an interface.
show mac access-group [**interface** *interface-i*d]	Displays MAC access lists applied to all Layer 2 interfaces or the specified Layer 2 interface.
show access-lists [*number* | *name*]	Displays the access list configuration.
show time-range	Verifies the time-range configuration.
show mac access-group [**interface** *interface-i*d]	Displays the MAC access list applied to the interface or all Layer 2 interfaces.

Configuration Examples for Network Security with ACLs

Creating a Standard ACL: Example

This example shows how to create a standard ACL to deny access to IP host 171.69.198.102, permit access to any others, and display the results.

```
Switch (config)# access-list 2 deny host 171.69.198.102
Switch (config)# access-list 2 permit any
Switch(config)# end
Switch# show access-lists
Standard IP access list 2
    10 deny   171.69.198.102
    20 permit any
```

Creating an Extended ACL: Example

This example shows how to create and display an extended access list to deny Telnet access from any host in network 171.69.198.0 to any host in network 172.20.52.0 and to permit any others. (The **eq** keyword after the destination address means to test for the TCP destination port number equaling Telnet.)

```
Switch(config)# access-list 102 deny tcp 171.69.198.0 0.0.0.255 172.20.52.0 0.0.0.255 eq
telnet
Switch(config)# access-list 102 permit tcp any any
```

```
Switch(config)# end
Switch# show access-lists
Extended IP access list 102
    10 deny tcp 171.69.198.0 0.0.0.255 172.20.52.0 0.0.0.255 eq telnet
    20 permit tcp any any
```

Configuring Time Ranges: Examples

This example shows how to configure time ranges for *workhours* and to configure January 1, 2006, as a company holiday and to verify your configuration.

```
Switch(config)# time-range workhours
Switch(config-time-range)# periodic weekdays 8:00 to 12:00
Switch(config-time-range)# periodic weekdays 13:00 to 17:00
Switch(config-time-range)# exit
Switch(config)# time-range new_year_day_2006
Switch(config-time-range)# absolute start 00:00 1 Jan 2006 end 23:59 1 Jan 2006
Switch(config-time-range)# end
Switch# show time-range
time-range entry: new_year_day_2003 (inactive)
   absolute start 00:00 01 January 2006 end 23:59 01 January 2006
time-range entry: workhours (inactive)
   periodic weekdays 8:00 to 12:00
   periodic weekdays 13:00 to 17:00
```

To apply a time range, enter the time-range name in an extended ACL that can implement time ranges. This example shows how to create and verify extended access list 188 that denies TCP traffic from any source to any destination during the defined holiday times and permits all TCP traffic during work hours.

```
Switch(config)# access-list 188 deny tcp any any time-range new_year_day_2006
Switch(config)# access-list 188 permit tcp any any time-range workhours
Switch(config)# end
Switch# show access-lists
Extended IP access list 188
   10 deny tcp any any time-range new_year_day_2006 (inactive)
   20 permit tcp any any time-range workhours (inactive)
```

Using Named ACLs: Example

This example uses named ACLs to permit and deny the same traffic.

```
Switch(config)# ip access-list extended deny_access
Switch(config-ext-nacl)# deny tcp any any time-range new_year_day_2006
Switch(config-ext-nacl)# exit
Switch(config)# ip access-list extended may_access
Switch(config-ext-nacl)# permit tcp any any time-range workhours
Switch(config-ext-nacl)# end
Switch# show ip access-lists
Extended IP access list lpip_default
    10 permit ip any any
Extended IP access list deny_access
    10 deny tcp any any time-range new_year_day_2006 (inactive)
Extended IP access list may_access
    10 permit tcp any any time-range workhours (inactive)
```

Including Comments in ACLs: Examples

In this example, the workstation that belongs to Jones is allowed access, and the workstation that belongs to Smith is not allowed access:

```
Switch(config)# access-list 1 remark Permit only Jones workstation through
Switch(config)# access-list 1 permit 171.69.2.88
Switch(config)# access-list 1 remark Do not allow Smith through
Switch(config)# access-list 1 deny 171.69.3.13
```

For an entry in a named IP ACL, use the **remark** access-list configuration command. To remove the remark, use the **no** form of this command.

In this example, the Jones subnet is not allowed to use outbound Telnet:

```
Switch(config)# ip access-list extended telnetting
Switch(config-ext-nacl)# remark Do not allow Jones subnet to telnet out
Switch(config-ext-nacl)# deny tcp host 171.69.2.88 any eq telnet
```

Applying ACL to a Port: Example

This example shows how to apply access list 2 to a port to filter packets entering the port:

```
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# ip access-group 2 in
```

Applying an ACL to an Interface: Example

For example, if you apply this ACL to an interface:

```
permit tcp source source-wildcard destination destination-wildcard range 5 60
permit tcp source source-wildcard destination destination-wildcard range 15 160
permit tcp source source-wildcard destination destination-wildcard range 115 1660
permit tcp source source-wildcard destination destination-wildcard
```

And if this message appears:

```
ACLMGR-2-NOVMR: Cannot generate hardware representation of access list [chars]
```

The flag-related operators are not available. To avoid this issue,

- Move the fourth ACE before the first ACE by using **ip access-list resequence** global configuration command:

  ```
  permit tcp source source-wildcard destination destination-wildcard
  permit tcp source source-wildcard destination destination-wildcard range 5 60
  permit tcp source source-wildcard destination destination-wildcard range 15 160
  permit tcp source source-wildcard destination destination-wildcard range 115 1660
  ```

  or

- Rename the ACL with a name or number that alphanumerically precedes the other ACLs (for example, rename ACL *79* to ACL *1*).

You can now apply the first ACE in the ACL to the interface. The switch allocates the ACE to available mapping bits in the Opselect index and then allocates flag-related operators to use the same bits in the TCAM.

Router ACLs function as follows:

- The hardware controls permit and deny actions of standard and extended ACLs (input and output) for security access control.
- If **log** has not been specified, the flows that match a *deny* statement in a security ACL are dropped by the hardware if *ip unreachables* is disabled. The flows matching a *permit* statement are switched in hardware.
- Adding the **log** keyword to an ACE in a router ACL causes a copy of the packet to be sent to the CPU for logging only. If the ACE is a *permit* statement, the packet is still switched and routed in hardware.

Routed ACLs: Examples

Figure 37-2 shows a small networked office environment with routed Port 2 connected to Server A, containing benefits and other information that all employees can access, and routed Port 1 connected to Server B, containing confidential payroll data. All users can access Server A, but Server B has restricted access.

Use router ACLs to do this in one of two ways:

- Create a standard ACL, and filter traffic coming to the server from Port 1.
- Create an extended ACL, and filter traffic coming from the server into Port 1.

Figure 37-2 Using Router ACLs to Control Traffic

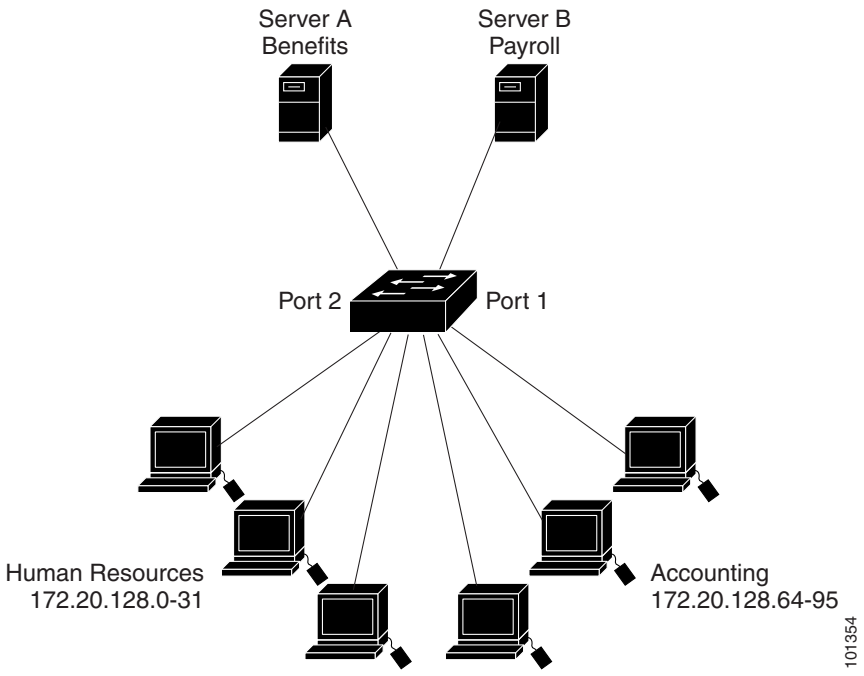

This example uses a standard ACL to filter traffic coming into Server B from a port, permitting traffic only from Accounting's source addresses 172.20.128.64 to 172.20.128.95. The ACL is applied to traffic coming out of routed Port 1 from the specified source address.

```
Switch(config)# access-list 6 permit 172.20.128.64 0.0.0.31
Switch(config)# end
```

```
Switch# show access-lists
Standard IP access list 6
 permit 172.20.128.64, wildcard bits 0.0.0.31
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# ip access-group 6 out
```

This example uses an extended ACL to filter traffic coming from Server B into a port, permitting traffic from any source address (in this case Server B) to only the Accounting destination addresses 172.20.128.64 to 172.20.128.95. The ACL is applied to traffic going into routed Port 1, permitting it to go only to the specified destination addresses. Note that with extended ACLs, you must enter the protocol (IP) before the source and destination information.

```
Switch(config)# access-list 106 permit ip any 172.20.128.64 0.0.0.31
Switch(config)# end
Switch# show access-lists
Extended IP access list 106
 permit ip any 172.20.128.64 0.0.0.31
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# ip access-group 106 in
```

Configuring Numbered ACLs: Example

In this example, network 36.0.0.0 is a Class A network whose second octet specifies a subnet; that is, its subnet mask is 255.255.0.0. The third and fourth octets of a network 36.0.0.0 address specify a particular host. Using access list 2, the switch accepts one address on subnet 48 and reject all others on that subnet. The last line of the list shows that the switch accepts addresses on all other network 36.0.0.0 subnets. The ACL is applied to packets entering a port.

```
Switch(config)# access-list 2 permit 36.48.0.3
Switch(config)# access-list 2 deny 36.48.0.0 0.0.255.255
Switch(config)# access-list 2 permit 36.0.0.0 0.255.255.255
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# ip access-group 2 in
```

Configuring Extended ACLs: Examples

In this example, the first line permits any incoming TCP connections with destination ports greater than 1023. The second line permits incoming TCP connections to the Simple Mail Transfer Protocol (SMTP) port of host 128.88.1.2. The third line permits incoming ICMP messages for error feedback.

```
Switch(config)# access-list 102 permit tcp any 128.88.0.0 0.0.255.255 gt 1023
Switch(config)# access-list 102 permit tcp any host 128.88.1.2 eq 25
Switch(config)# access-list 102 permit icmp any any
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# ip access-group 102 in
```

In this example, suppose that you have a network connected to the Internet, and you want any host on the network to be able to form TCP connections to any host on the Internet. However, you do not want IP hosts to be able to form TCP connections to hosts on your network, except to the mail (SMTP) port of a dedicated mail host.

SMTP uses TCP port 25 on one end of the connection and a random port number on the other end. The same port numbers are used throughout the life of the connection. Mail packets coming in from the Internet have a destination port of 25. Outbound packets have the port numbers reversed. Because the secure system of the network always accepts mail connections on port 25, the incoming and outgoing services are separately controlled. The ACL must be configured as an input ACL on the outbound interface and an output ACL on the inbound interface.

In this example, the network is a Class B network with the address 128.88.0.0, and the mail host address is 128.88.1.2. The **established** keyword is used only for the TCP to show an established connection. A match occurs if the TCP datagram has the ACK or RST bits set, which show that the packet belongs to an existing connection. Gigabit Ethernet interface 1 is the interface that connects the router to the Internet.

```
Switch(config)# access-list 102 permit tcp any 128.88.0.0 0.0.255.255 established
Switch(config)# access-list 102 permit tcp any host 128.88.1.2 eq 25
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# ip access-group 102 in
```

Creating Named ACLs: Example

This example creates a standard ACL named *Internet_filter* and an extended ACL named *marketing_group*. The *Internet_filter* ACL allows all traffic from the source address 1.2.3.4.

```
Switch(config)# ip access-list standard Internet_filter
Switch(config-ext-nacl)# permit 1.2.3.4
Switch(config-ext-nacl)# exit
```

The *marketing_group* ACL allows any TCP Telnet traffic to the destination address and wildcard 171.69.0.0 0.0.255.255 and denies any other TCP traffic. It permits ICMP traffic, denies UDP traffic from any source to the destination address range 171.69.0.0 through 179.69.255.255 with a destination port less than 1024, denies any other IP traffic, and provides a log of the result.

```
Switch(config)# ip access-list extended marketing_group
Switch(config-ext-nacl)# permit tcp any 171.69.0.0 0.0.255.255 eq telnet
Switch(config-ext-nacl)# deny tcp any any
Switch(config-ext-nacl)# permit icmp any any
Switch(config-ext-nacl)# deny udp any 171.69.0.0 0.0.255.255 lt 1024
Switch(config-ext-nacl)# deny ip any any log
Switch(config-ext-nacl)# exit
```

The *Internet_filter* ACL is applied to outgoing traffic and the *marketing_group* ACL is applied to incoming traffic on a Layer 3 port.

```
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# no switchport
Switch(config-if)# ip address 2.0.5.1 255.255.255.0
Switch(config-if)# ip access-group Internet_filter out
Switch(config-if)# ip access-group marketing_group in
```

Applying Time Range to an IP ACL: Example

This example denies HTTP traffic on IP on Monday through Friday between the hours of 8:00 a.m. and 6:00 p.m (18:00). The example allows UDP traffic only on Saturday and Sunday from noon to 8:00 p.m. (20:00).

```
Switch(config)# time-range no-http
Switch(config)# periodic weekdays 8:00 to 18:00
!
Switch(config)# time-range udp-yes
Switch(config)# periodic weekend 12:00 to 20:00
!
Switch(config)# ip access-list extended strict
Switch(config-ext-nacl)# deny tcp any any eq www time-range no-http
Switch(config-ext-nacl)# permit udp any any time-range udp-yes
!
Switch(config-ext-nacl)# exit
Switch(config)# interface gigabitethernet1/1
```

```
Switch(config-if)# ip access-group strict in
```

Creating Commented IP ACL Entries: Examples

In this example of a numbered ACL, the workstation that belongs to Jones is allowed access, and the workstation that belongs to Smith is not allowed access:

```
Switch(config)# access-list 1 remark Permit only Jones workstation through
Switch(config)# access-list 1 permit 171.69.2.88
Switch(config)# access-list 1 remark Do not allow Smith workstation through
Switch(config)# access-list 1 deny 171.69.3.13
```

In this example of a numbered ACL, the Winter and Smith workstations are not allowed to browse the web:

```
Switch(config)# access-list 100 remark Do not allow Winter to browse the web
Switch(config)# access-list 100 deny host 171.69.3.85 any eq www
Switch(config)# access-list 100 remark Do not allow Smith to browse the web
Switch(config)# access-list 100 deny host 171.69.3.13 any eq www
```

In this example of a named ACL, the Jones subnet is not allowed access:

```
Switch(config)# ip access-list standard prevention
Switch(config-std-nacl)# remark Do not allow Jones subnet through
Switch(config-std-nacl)# deny 171.69.0.0 0.0.255.255
```

In this example of a named ACL, the Jones subnet is not allowed to use outbound Telnet:

```
Switch(config)# ip access-list extended telnetting
Switch(config-ext-nacl)# remark Do not allow Jones subnet to telnet out
Switch(config-ext-nacl)# deny tcp 171.69.0.0 0.0.255.255 any eq telnet
```

Configuring ACL Logging: Examples

Two variations of logging are supported on router ACLs. The **log** keyword sends an informational logging message to the console about the packet that matches the entry; the **log-input** keyword includes the input interface in the log entry.

In this example, standard named access list *stan1* denies traffic from 10.1.1.0 0.0.0.255, allows traffic from all other sources, and includes the **log** keyword.

```
Switch(config)# ip access-list standard stan1
Switch(config-std-nacl)# deny 10.1.1.0 0.0.0.255 log
Switch(config-std-nacl)# permit any log
Switch(config-std-nacl)# exit
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# ip access-group stan1 in
Switch(config-if)# end
Switch# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Console logging: level debugging, 37 messages logged
    Monitor logging: level debugging, 0 messages logged
    Buffer logging: level debugging, 37 messages logged
    File logging: disabled
    Trap logging: level debugging, 39 message lines logged

Log Buffer (4096 bytes):

00:00:48: NTP: authentication delay calculation problems
```

```
<output truncated>

00:09:34:%SEC-6-IPACCESSLOGS:list stan1 permitted 0.0.0.0 1 packet
00:09:59:%SEC-6-IPACCESSLOGS:list stan1 denied 10.1.1.15 1 packet
00:10:11:%SEC-6-IPACCESSLOGS:list stan1 permitted 0.0.0.0 1 packet
```

This example is a named extended access list *ext1* that permits ICMP packets from any source to 10.1.1.0 0.0.0.255 and denies all UDP packets.

```
Switch(config)# ip access-list extended ext1
Switch(config-ext-nacl)# permit icmp any 10.1.1.0 0.0.0.255 log
Switch(config-ext-nacl)# deny udp any any log
Switch(config-std-nacl)# exit
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# ip access-group ext1 in
```

Applying a MAC ACL to a Layer 2 Interface: Examples

This example shows how to create and display an access list named *mac1*, denying only EtherType DECnet Phase IV traffic, but permitting all other types of traffic.

```
Switch(config)# mac access-list extended mac1
Switch(config-ext-macl)# deny any any decnet-iv
Switch(config-ext-macl)# permit any any
Switch(config-ext-macl)# end
Switch # show access-lists
Extended MAC access list mac1
    10 deny   any any decnet-iv
    20 permit any any
```

This example shows how to apply MAC access list *mac1* to a port to filter packets entering the port:

```
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# mac access-group mac1 in
```

Note The **mac access-group** interface configuration command is only valid when applied to a physical Layer 2 interface.You cannot use the command on EtherChannel port channels.

After receiving a packet, the switch checks it against the inbound ACL. If the ACL permits it, the switch continues to process the packet. If the ACL rejects the packet, the switch discards it. When you apply an undefined ACL to an interface, the switch acts as if the ACL has not been applied and permits all packets. Remember this behavior if you use undefined ACLs for network security.

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
Cisco IOS multicast commands	*Cisco IOS IP Command Reference, Volume 3 of 3:Multicast*
Cisco IOS IP Addressing and Services configuration	*Cisco IOS IP Configuration Guide*
Cisco IOS ACL configuration	*Cisco IOS IP Command Reference, Volume 1 of 3: Addressing and Services* *Cisco IOS Security Configuration Guide*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 38

Configuring Standard QoS

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for Standard QoS

Before configuring standard QoS, you must have a thorough understanding of these items:

- The types of applications used and the traffic patterns on your network.
- Traffic characteristics and needs of your network. Is the traffic bursty? Do you need to reserve bandwidth for voice and video streams?
- Bandwidth requirements and speed of the network.
- Location of congestion points in the network.

Restrictions for Standard QoS

- To use this feature, the switch must be running the LAN Base image.
- Control traffic (such as spanning-tree bridge protocol data units [BPDUs] and routing update packets) received by the switch are subject to all ingress QoS processing.
- You are likely to lose data when you change queue settings; therefore, try to make changes when traffic is at a minimum.
- The IPv6 QoS trust feature is not supported.

Information About Standard QoS

This chapter describes how to configure quality of service (QoS) by using automatic QoS (auto-QoS) commands or by using standard QoS commands on the switch. With QoS, you can provide preferential treatment to certain types of traffic at the expense of others. Without QoS, the switch offers best-effort service to each packet, regardless of the packet contents or size. It sends the packets without any assurance of reliability, delay bounds, or throughput.

You can configure QoS on physical ports and on switch virtual interfaces (SVIs). Other than to apply policy maps, you configure the QoS settings, such as classification, queueing, and scheduling, the same way on physical ports and SVIs. When configuring QoS on a physical port, you apply a nonhierarchical policy map to a port. When configuring QoS on an SVI, you apply a nonhierarchical or a hierarchical policy map.

The switch supports some of the modular QoS CLI (MQC) commands. For more information about the MQC commands, see the "Modular Quality of Service Command-Line Interface Overview" chapter of the *Cisco IOS Quality of Service Solutions Guide, Release 12.2*.

Typically, networks operate on a best-effort delivery basis, which means that all traffic has equal priority and an equal chance of being delivered in a timely manner. When congestion occurs, all traffic has an equal chance of being dropped.

When you configure the QoS feature, you can select specific network traffic, prioritize it according to its relative importance, and use congestion-management and congestion-avoidance techniques to provide preferential treatment. Implementing QoS in your network makes network performance more predictable and bandwidth utilization more effective.

The QoS implementation is based on the Differentiated Services (Diff-Serv) architecture, an emerging standard from the Internet Engineering Task Force (IETF). This architecture specifies that each packet is classified upon entry into the network.

The classification is carried in the IP packet header, using 6 bits from the deprecated IP type of service (ToS) field to carry the classification (*class*) information. Classification can also be carried in the Layer 2 frame. These special bits in the Layer 2 frame or a Layer 3 packet are described here and shown in Figure 38-1:

- Prioritization bits in Layer 2 frames:

  Layer 2 IEEE 802.1Q frame headers have a 2-byte Tag Control Information field that carries the CoS value in the three most-significant bits, which are called the User Priority bits. On ports configured as Layer 2 IEEE 802.1Q trunks, all traffic is in IEEE 802.1Q frames except for traffic in the native VLAN.

  Other frame types cannot carry Layer 2 CoS values.

  Layer 2 CoS values range from 0 for low priority to 7 for high priority.

- Prioritization bits in Layer 3 packets:

  Layer 3 IP packets can carry either an IP precedence value or a Differentiated Services Code Point (DSCP) value. QoS supports the use of either value because DSCP values are backward-compatible with IP precedence values.

  IP precedence values range from 0 to 7.

  DSCP values range from 0 to 63.

Note IPv6 port-based trust with the dual IPv4 and IPv6 Switch Database Management (SDM) templates is supported on this switch. You must reload the switch with the dual IPv4 and IPv6 templates for switches running IPv6. For more information, see Chapter 11, "Configuring SDM Templates."

Figure 38-1 QoS Classification Layers in Frames and Packets

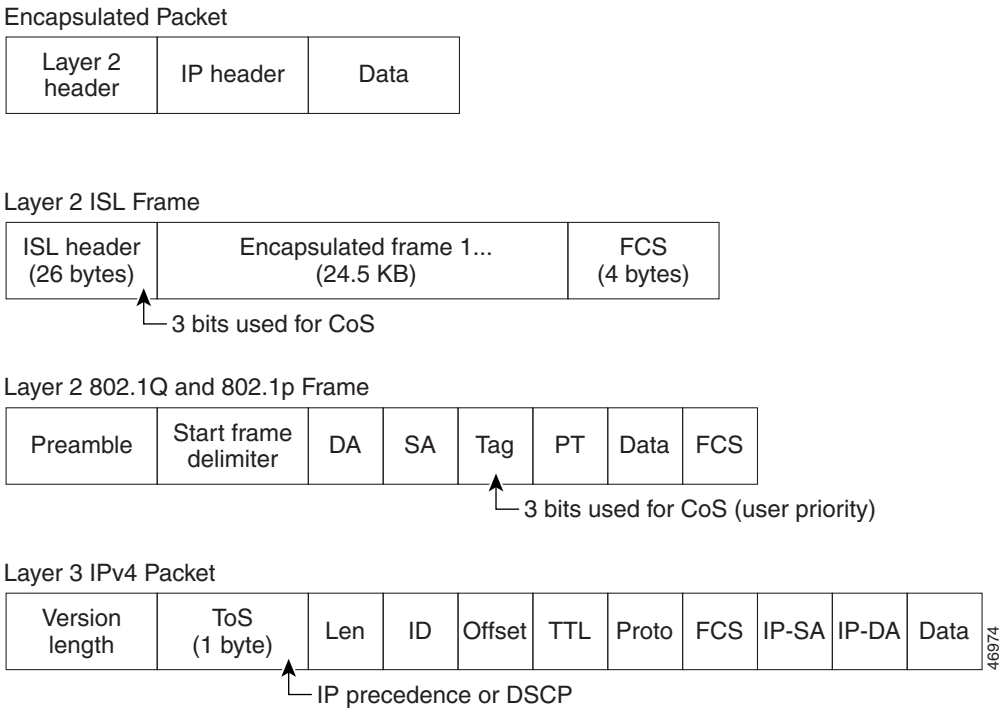

All switches and routers that access the Internet rely on the class information to provide the same forwarding treatment to packets with the same class information and different treatment to packets with different class information. The class information in the packet can be assigned by end hosts or by switches or routers along the way, based on a configured policy, detailed examination of the packet, or both. Detailed examination of the packet is expected to happen closer to the edge of the network so that the core switches and routers are not overloaded with this task.

Switches and routers along the path can use the class information to limit the amount of resources allocated per traffic class. The behavior of an individual device when handling traffic in the DiffServ architecture is called per-hop behavior. If all devices along a path provide a consistent per-hop behavior, you can construct an end-to-end QoS solution.

Implementing QoS in your network can be a simple or complex task and depends on the QoS features offered by your internetworking devices, the traffic types and patterns in your network, and the granularity of control that you need over incoming and outgoing traffic.

Standard QoS Model

To implement QoS, the switch must distinguish packets or flow from one another (classify), assign a label to indicate the given quality of service as the packets move through the switch, make the packets comply with the configured resource usage limits (police and mark), and provide different treatment (queue and schedule) in all situations where resource contention exists. The switch also needs to ensure that traffic sent from it meets a specific traffic profile (shape).

Figure 38-2 shows the standard QoS model. Actions at the ingress port include classifying traffic, policing, marking, queueing, and scheduling:

- Classifying a distinct path for a packet by associating it with a QoS label. The switch maps the CoS or DSCP in the packet to a QoS label to distinguish one kind of traffic from another. The QoS label that is generated identifies all future QoS actions to be performed on this packet. For more information, see the "Classification" section on page 38-10.
- Policing determines whether a packet is in or out of profile by comparing the rate of the incoming traffic to the configured policer. The policer limits the bandwidth consumed by a flow of traffic. The result is passed to the marker. For more information, see the "Policing and Marking" section on page 38-14.
- Marking evaluates the policer and configuration information for the action to be taken when a packet is out of profile and determines what to do with the packet (pass through a packet without modification, mark down the QoS label in the packet, or drop the packet). For more information, see the "Policing and Marking" section on page 38-14.
- Queueing evaluates the QoS label and the corresponding DSCP or CoS value to select into which of the two ingress queues to place a packet. Queueing is enhanced with the weighted tail-drop (WTD) algorithm, a congestion-avoidance mechanism. If the threshold is exceeded, the packet is dropped. For more information, see the "Queueing and Scheduling Overview" section on page 38-19.
- Scheduling services the queues based on their configured shaped round robin (SRR) weights. One of the ingress queues is the priority queue, and SRR services it for its configured share before servicing the other queue. For more information, see the "SRR Shaping and Sharing" section on page 38-20.

Actions at the egress port include queueing and scheduling:

- Queueing evaluates the QoS packet label and the corresponding DSCP or CoS value before selecting which of the four egress queues to use. Because congestion can occur when multiple ingress ports simultaneously send data to an egress port, WTD differentiates traffic classes and subjects the packets to different thresholds based on the QoS label. If the threshold is exceeded, the packet is dropped. For more information, see the "Queueing and Scheduling Overview" section on page 38-19.
- Scheduling services the four egress queues based on their configured SRR shared or shaped weights. One of the queues (queue 1) can be the expedited queue, which is serviced until empty before the other queues are serviced.

Figure 38-2 Standard QoS Model

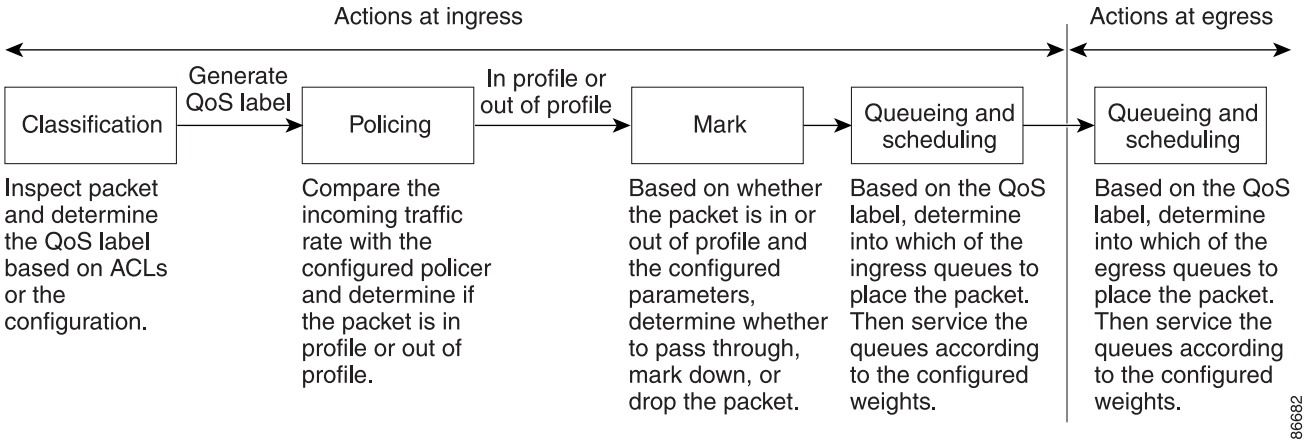

Standard QoS Configuration Guidelines

QoS ACL

These are the guidelines for configuring QoS with access control lists (ACLs):

- It is not possible to match IP fragments against configured IP extended ACLs to enforce QoS. IP fragments are sent as best-effort. IP fragments are denoted by fields in the IP header.
- Only one ACL per class map and only one **match** class-map configuration command per class map are supported. The ACL can have multiple ACEs, which match fields against the contents of the packet.
- A trust statement in a policy map requires multiple TCAM entries per ACL line. If an input service policy map contains a trust statement in an ACL, the access-list might be too large to fit into the available QoS TCAM and an error can occur when you apply the policy map to a port. Whenever possible, you should minimize the number of lines in a QoS ACL.

QoS on Interfaces

These are the guidelines for configuring QoS on physical ports. This section also applies to SVIs (Layer 3 interfaces):

- You can configure QoS on physical ports and SVIs. When configuring QoS on physical ports, you create and apply nonhierarchical policy maps. When configuring QoS on SVIs, you can create and apply nonhierarchical and hierarchical policy maps.
- Incoming traffic is classified, policed, and marked down (if configured) regardless of whether the traffic is bridged, routed, or sent to the CPU. It is possible for bridged frames to be dropped or to have their DSCP and CoS values modified.
- Follow these guidelines when configuring policy maps on physical ports or SVIs:
  - You cannot apply the same policy map to a physical port and to an SVI.
  - If VLAN-based QoS is configured on a physical port, the switch removes all the port-based policy maps on the port. The traffic on this physical port is now affected by the policy map attached to the SVI to which the physical port belongs.

- In a hierarchical policy map attached to an SVI, you can only configure an individual policer at the interface level on a physical port to specify the bandwidth limits for the traffic on the port. The ingress port must be configured as a trunk or as a static-access port. You cannot configure policers at the VLAN level of the hierarchical policy map.
- The switch does not support aggregate policers in hierarchical policy maps.
- After the hierarchical policy map is attached to an SVI, the interface-level policy map cannot be modified or removed from the hierarchical policy map. A new interface-level policy map also cannot be added to the hierarchical policy map. If you want these changes to occur, the hierarchical policy map must first be removed from the SVI. You also cannot add or remove a class map specified in the hierarchical policy map.

Policing

- The port ASIC device, which controls more than one physical port, supports 256 policers (255 user-configurable policers plus 1 policer reserved for system internal use). The maximum number of user-configurable policers supported per port is 63. For example, you could configure 32 policers on a Gigabit Ethernet port and 8 policers on a Fast Ethernet port, or you could configure 64 policers on a Gigabit Ethernet port and 5 policers on a Fast Ethernet port. Policers are allocated on demand by the software and are constrained by the hardware and ASIC boundaries. You cannot reserve policers per port; there is no guarantee that a port will be assigned to any policer.
- Only one policer is applied to a packet on an ingress port. Only the average rate and committed burst parameters are configurable.
- You can create an aggregate policer that is shared by multiple traffic classes within the same nonhierarchical policy map. However, you cannot use the aggregate policer across different policy maps.
- On a port configured for QoS, all traffic received through the port is classified, policed, and marked according to the policy map attached to the port. On a trunk port configured for QoS, traffic in *all* VLANs received through the port is classified, policed, and marked according to the policy map attached to the port.
- If you have EtherChannel ports configured on your switch, you must configure QoS classification, policing, mapping, and queueing on the individual physical ports that comprise the EtherChannel. You must decide whether the QoS configuration should match on all ports in the EtherChannel.

Default Standard QoS Configuration

QoS is disabled. There is no concept of trusted or untrusted ports because the packets are not modified (the CoS, DSCP, and IP precedence values in the packet are not changed). Traffic is switched in pass-through mode (packets are switched without any rewrites and classified as best effort without any policing).

When QoS is enabled with the **mls qos** global configuration command and all other QoS settings are at their defaults, traffic is classified as best effort (the DSCP and CoS value is set to 0) without any policing. No policy maps are configured. The default port trust state on all ports is untrusted. The default ingress and egress queue settings are described in the "Default Ingress Queue Settings" section on page 38-7 and the "Default Egress Queue Settings" section on page 38-7.

Default Ingress Queue Settings

Table 38-1 shows the default ingress queue settings when QoS is enabled.

Table 38-1 Default Ingress Queue Settings

Feature	Queue 1	Queue 2
Buffer allocation	90 percent	10 percent
Bandwidth allocation [1]	4	4
Priority queue bandwidth [2]	0	10
WTD drop threshold 1	100 percent	100 percent
WTD drop threshold 2	100 percent	100 percent

1. The bandwidth is equally shared between the queues. SRR sends packets in shared mode only.
2. Queue 2 is the priority queue. SRR services the priority queue for its configured share before servicing the other queue.

Table 38-2 shows the default CoS input queue threshold map when QoS is enabled.

Table 38-2 Default CoS Input Queue Threshold Map

CoS Value	Queue ID–Threshold ID
0–4	1–1
5	2–1
6, 7	1–1

Table 38-3 shows the default DSCP input queue threshold map when QoS is enabled.

Table 38-3 Default DSCP Input Queue Threshold Map

DSCP Value	Queue ID–Threshold ID
0–39	1–1
40–47	2–1
48–63	1–1

Default Egress Queue Settings

Table 38-4 shows the default egress queue settings for each queue-set when QoS is enabled. All ports are mapped to queue-set 1. The port bandwidth limit is set to 100 percent and rate unlimited.

Table 38-4 Default Egress Queue Settings

Feature	Queue 1	Queue 2	Queue 3	Queue 4
Buffer allocation	25 percent	25 percent	25 percent	25 percent
WTD drop threshold 1	100 percent	200 percent	100 percent	100 percent
WTD drop threshold 2	100 percent	200 percent	100 percent	100 percent
Reserved threshold	50 percent	50 percent	50 percent	50 percent

Table 38-4 Default Egress Queue Settings (continued)

Feature	Queue 1	Queue 2	Queue 3	Queue 4
Maximum threshold	400 percent	400 percent	400 percent	400 percent
SRR shaped weights (absolute) [1]	25	0	0	0
SRR shared weights [2]	25	25	25	25

1. A shaped weight of zero means that this queue is operating in shared mode.
2. One quarter of the bandwidth is allocated to each queue.

Table 38-5 shows the default CoS output queue threshold map when QoS is enabled.

Table 38-5 Default CoS Output Queue Threshold Map

CoS Value	Queue ID–Threshold ID
0, 1	2–1
2, 3	3–1
4	4–1
5	1–1
6, 7	4–1

Table 38-6 shows the default DSCP output queue threshold map when QoS is enabled.

Table 38-6 Default DSCP Output Queue Threshold Map

DSCP Value	Queue ID–Threshold ID
0–15	2–1
16–31	3–1
32–39	4–1
40–47	1–1
48–63	4–1

Default Mapping Table Settings

Note If these values are not appropriate for your network, you need to modify them.

Table 38-7 shows the DSCP-to-CoS map to generate a CoS value, which is used to select one of the four egress queues.

Table 38-7 Default DSCP-to-CoS Map

DSCP Value	CoS Value
0–7	0
8–15	1

Table 38-7 Default DSCP-to-CoS Map (continued)

DSCP Value	CoS Value
16–23	2
24–31	3
32–39	4
40–47	5
48–55	6
56–63	7

Table 38-8 shows the IP-precedence-to-DSCP map to map IP precedence values in incoming packets to a DSCP value that QoS uses internally to represent the priority of the traffic.

Table 38-8 Default IP-Precedence-to-DSCP Map

IP Precedence Value	DSCP Value
0	0
1	8
2	16
3	24
4	32
5	40
6	48
7	56

Table 38-9 shows the CoS-to-DSCP map to map CoS values in incoming packets to a DSCP value that QoS uses internally to represent the priority of the traffic.

Table 38-9 CoS-to-DSCP Map

CoS Value	DSCP Value
0	0
1	8
2	16
3	24
4	32
5	40
6	48
7	56

The default DSCP-to-DSCP-mutation map is a null map, which maps an incoming DSCP value to the same DSCP value.

The default policed-DSCP map is a null map, which maps an incoming DSCP value to the same DSCP value (no markdown).

Classification

Classification is the process of distinguishing one kind of traffic from another by examining the fields in the packet. Classification is enabled only if QoS is globally enabled on the switch. By default, QoS is globally disabled, so no classification occurs.

During classification, the switch performs a lookup and assigns a QoS label to the packet. The QoS label identifies all QoS actions to be performed on the packet and from which queue the packet is sent.

The QoS label is based on the DSCP or the CoS value in the packet and decides the queueing and scheduling actions to perform on the packet. The label is mapped according to the trust setting and the packet type as shown in Figure 38-3 on page 38-12.

You specify which fields in the frame or packet that you want to use to classify incoming traffic. For non-IP traffic, you have these classification options as shown in Figure 38-3:

- Trust the CoS value in the incoming frame (configure the port to trust CoS). Then use the configurable CoS-to-DSCP map to generate a DSCP value for the packet. Layer 2 ISL frame headers carry the CoS value in the 3 least-significant bits of the 1-byte User field. Layer 2 IEEE 802.1Q frame headers carry the CoS value in the 3 most-significant bits of the Tag Control Information field. CoS values range from 0 for low priority to 7 for high priority.
- Trust the DSCP or trust IP precedence value in the incoming frame. These configurations are meaningless for non-IP traffic. If you configure a port with either of these options and non-IP traffic is received, the switch assigns a CoS value and generates an internal DSCP value from the CoS-to-DSCP map. The switch uses the internal DSCP value to generate a CoS value representing the priority of the traffic.
- Perform the classification based on a configured Layer 2 MAC access control list (ACL), which can examine the MAC source address, the MAC destination address, and other fields. If no ACL is configured, the packet is assigned 0 as the DSCP and CoS values, which means best-effort traffic. Otherwise, the policy-map action specifies a DSCP or CoS value to assign to the incoming frame.

For IP traffic, you have these classification options as shown in Figure 38-3:

- Trust the DSCP value in the incoming packet (configure the port to trust DSCP), and assign the same DSCP value to the packet. The IETF defines the 6 most-significant bits of the 1-byte ToS field as the DSCP. The priority represented by a particular DSCP value is configurable. DSCP values range from 0 to 63.

  For ports that are on the boundary between two QoS administrative domains, you can modify the DSCP to another value by using the configurable DSCP-to-DSCP-mutation map.
- Trust the IP precedence value in the incoming packet (configure the port to trust IP precedence), and generate a DSCP value for the packet by using the configurable IP-precedence-to-DSCP map. The IP Version 4 specification defines the 3 most-significant bits of the 1-byte ToS field as the IP precedence. IP precedence values range from 0 for low priority to 7 for high priority.
- Trust the CoS value (if present) in the incoming packet, and generate a DSCP value for the packet by using the CoS-to-DSCP map. If the CoS value is not present, use the default port CoS value.

- Perform the classification based on a configured IP standard or an extended ACL, which examines various fields in the IP header. If no ACL is configured, the packet is assigned 0 as the DSCP and CoS values, which means best-effort traffic. Otherwise, the policy-map action specifies a DSCP or CoS value to assign to the incoming frame.

For information on the maps described in this section, see the "Mapping Tables" section on page 38-18. For configuration information on port trust states, see the "Configuring Classification Using Port Trust States" section on page 38-32.

After classification, the packet is sent to the policing, marking, and the ingress queueing and scheduling stages.

Figure 38-3 Classification Flowchart

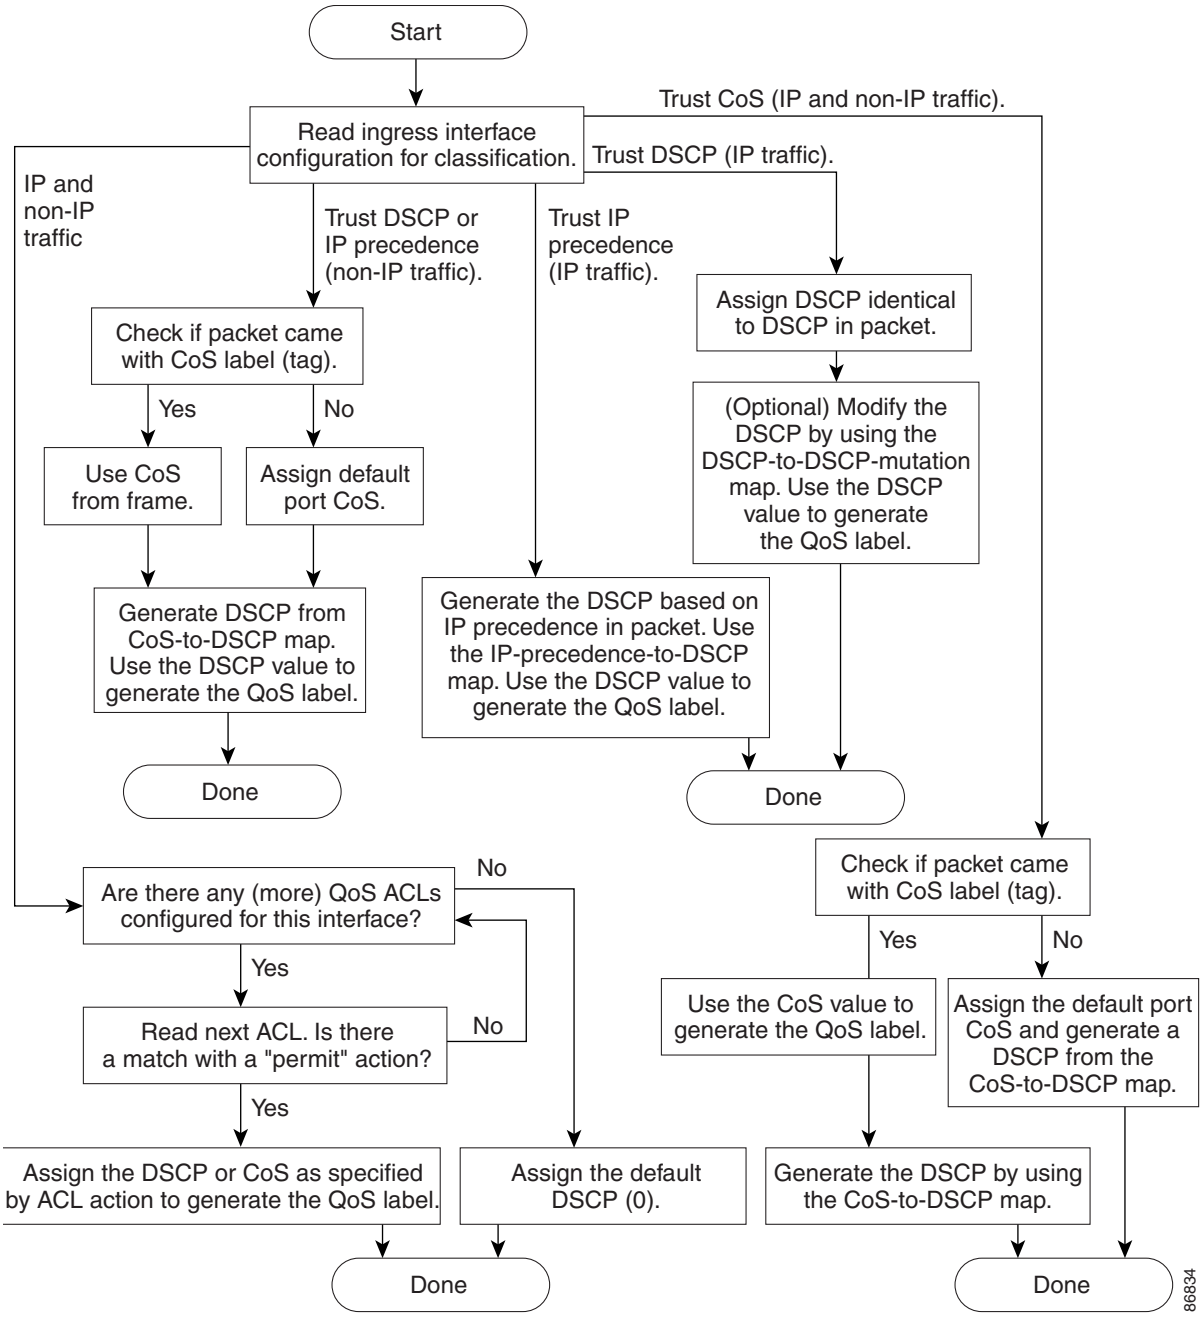

Classification Based on QoS ACLs

You can use IP standard, IP extended, or Layer 2 MAC ACLs to define a group of packets with the same characteristics (*class*). In the QoS context, the permit and deny actions in the access control entries (ACEs) have different meanings than with security ACLs:

- If a match with a permit action is encountered (first-match principle), the specified QoS-related action is taken.
- If a match with a deny action is encountered, the ACL being processed is skipped, and the next ACL is processed.
- If no match with a permit action is encountered and all the ACEs have been examined, no QoS processing occurs on the packet, and the switch offers best-effort service to the packet.
- If multiple ACLs are configured on a port, the lookup stops after the packet matches the first ACL with a permit action, and QoS processing begins.

Note When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.

After a traffic class has been defined with the ACL, you can attach a policy to it. A policy might contain multiple classes with actions specified for each one of them. A policy might include commands to classify the class as a particular aggregate (for example, assign a DSCP) or rate-limit the class. This policy is then attached to a particular port on which it becomes effective.

You implement IP ACLs to classify IP traffic by using the **access-list** global configuration command; you implement Layer 2 MAC ACLs to classify non-IP traffic by using the **mac access-list extended** global configuration command. For configuration information, see the "Configuring a QoS Policy" section on page 38-36.

Classification Based on Class Maps and Policy Maps

A class map is a mechanism that you use to name a specific traffic flow (or class) and to isolate it from all other traffic. The class map defines the criteria used to match against a specific traffic flow to further classify it. The criteria can include matching the access group defined by the ACL or matching a specific list of DSCP or IP precedence values. If you have more than one type of traffic that you want to classify, you can create another class map and use a different name. After a packet is matched against the class-map criteria, you further classify it through the use of a policy map.

A policy map specifies which traffic class to act on. Actions can include trusting the CoS, DSCP, or IP precedence values in the traffic class; setting a specific DSCP or IP precedence value in the traffic class; or specifying the traffic bandwidth limitations and the action to take when the traffic is out of profile. Before a policy map can be effective, you must attach it to a port.

You create a class map by using the **class-map** global configuration command or the **class** policy-map configuration command. You should use the **class-map** command when the map is shared among many ports. When you enter the **class-map** command, the switch enters the class-map configuration mode. In this mode, you define the match criterion for the traffic by using the **match** class-map configuration command.

You create and name a policy map by using the **policy-map** global configuration command. When you enter this command, the switch enters the policy-map configuration mode. In this mode, you specify the actions to take on a specific traffic class by using the **class**, **trust**, or **set** policy-map configuration and policy-map class configuration commands.

The policy map can contain the **police** and **police aggregate** policy-map class configuration commands, which define the policer, the bandwidth limitations of the traffic, and the action to take if the limits are exceeded.

To enable the policy map, you attach it to a port by using the **service-policy** interface configuration command.

You can apply a nonhierarchical policy map to a physical port or an SVI. However, a hierarchical policy map can only be applied to an SVI. A hierarchical policy map contains two levels. The first level, the VLAN level, specifies the actions to be taken against a traffic flow on the SVI. The second level, the interface level, specifies the actions to be taken against the traffic on the physical ports that belong to the SVI. The interface-level actions are specified in the interface-level policy map.

For more information, see the "Policing and Marking" section on page 38-14. For configuration information, see the "Configuring a QoS Policy" section on page 38-36.

Policing and Marking

After a packet is classified and has a DSCP-based or CoS-based QoS label assigned to it, the policing and marking process can begin as shown in Figure 38-4.

Policing involves creating a policer that specifies the bandwidth limits for the traffic. Packets that exceed the limits are *out of profile* or *nonconforming*. Each policer decides on a packet-by-packet basis whether the packet is in or out of profile and specifies the actions on the packet. These actions, carried out by the marker, include passing through the packet without modification, dropping the packet, or modifying (marking down) the assigned DSCP of the packet and allowing the packet to pass through. The configurable policed-DSCP map provides the packet with a new DSCP-based QoS label. For information on the policed-DSCP map, see the "Mapping Tables" section on page 38-18. Marked-down packets use the same queues as the original QoS label to prevent packets in a flow from getting out of order.

Note All traffic, regardless of whether it is bridged or routed, is subjected to a policer, if one is configured. As a result, bridged packets might be dropped or might have their DSCP or CoS fields modified when they are policed and marked.

You can configure policing on a physical port or an SVI. On a physical port, you can configure the trust state, set a new DSCP or IP precedence value in the packet, or define an individual or aggregate policer. For more information about configuring policing on physical ports, see the "Policing on Physical Ports" section on page 38-15. When configuring policy maps on an SVI, you can create a hierarchical policy map and can define an individual policer only in the secondary interface-level policy map. For more information, see the "Policing on SVIs" section on page 38-16.

After you configure the policy map and policing actions, attach the policy to an ingress port or SVI by using the **service-policy** interface configuration command.

Policing on Physical Ports

In policy maps on physical ports, you can create these types of policers:

- Individual—QoS applies the bandwidth limits specified in the policer separately to each matched traffic class. You configure this type of policer within a policy map by using the **police** policy-map class configuration command.
- Aggregate—QoS applies the bandwidth limits specified in an aggregate policer cumulatively to all matched traffic flows. You configure this type of policer by specifying the aggregate policer name within a policy map by using the **police aggregate** policy-map class configuration command. You specify the bandwidth limits of the policer by using the **mls qos aggregate-policer** global configuration command. In this way, the aggregate policer is shared by multiple classes of traffic within a policy map.

Note You can only configure individual policers on an SVI.

Policing uses a token-bucket algorithm. As each frame is received by the switch, a token is added to the bucket. The bucket has a hole in it and leaks at a rate that you specify as the average traffic rate in bits per second. Each time a token is added to the bucket, the switch verifies that there is enough room in the bucket. If there is not enough room, the packet is marked as nonconforming, and the specified policer action is taken (dropped or marked down).

How quickly the bucket fills is a function of the bucket depth (burst-byte), the rate at which the tokens are removed (rate-b/s), and the duration of the burst above the average rate. The size of the bucket imposes an upper limit on the burst length and limits the number of frames that can be transmitted back-to-back. If the burst is short, the bucket does not overflow, and no action is taken against the traffic flow. However, if a burst is long and at a higher rate, the bucket overflows, and the policing actions are taken against the frames in that burst.

You configure the bucket depth (the maximum burst that is tolerated before the bucket overflows) by using the *burst-byte* option of the **police** policy-map class configuration command or the **mls qos aggregate-policer** global configuration command. You configure how fast (the average rate) that the tokens are removed from the bucket by using the *rate-bps* option of the **police** policy-map class configuration command or the **mls qos aggregate-policer** global configuration command.

Figure 38-4 shows the policing and marking process. These types of policy maps are configured:

- A nonhierarchical policy map on a physical port.
- The interface level of a hierarchical policy map attached to an SVI. The physical ports are specified in this secondary policy map.

Figure 38-4 Policing and Marking Flowchart on Physical Ports

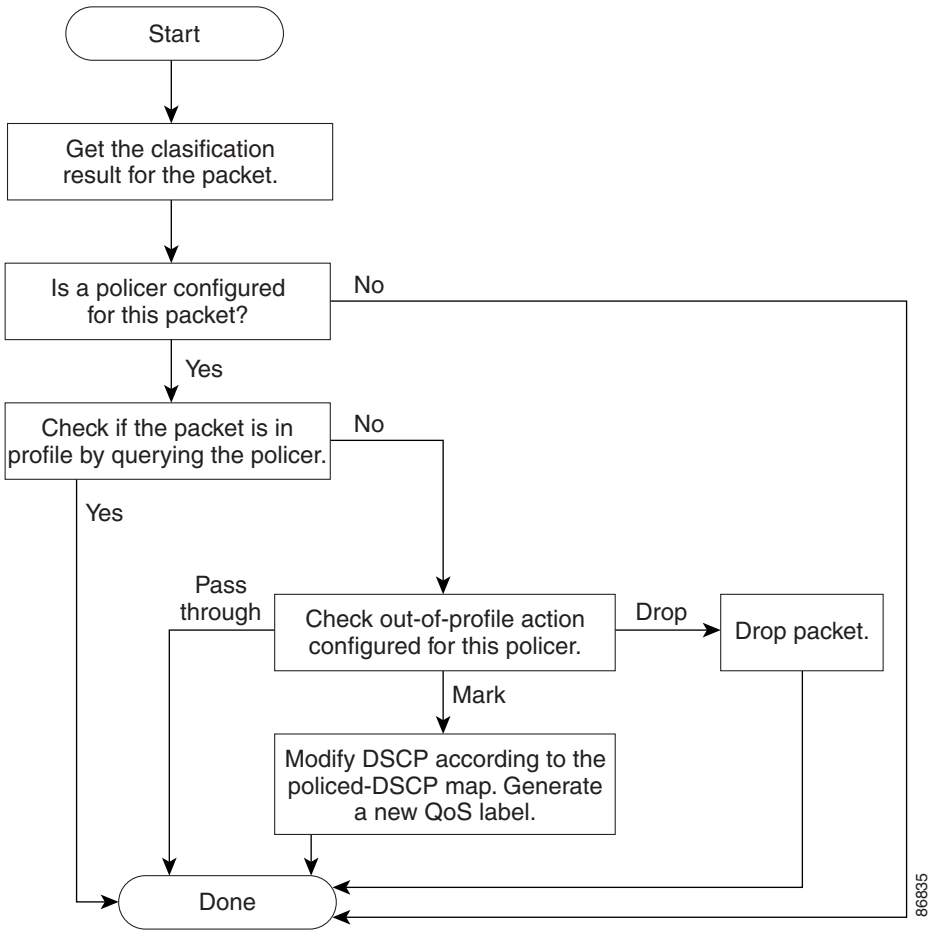

Policing on SVIs

Note Before configuring a hierarchical policy map with individual policers on an SVI, you must enable VLAN-based QoS on the physical ports that belong to the SVI. Though a policy map is attached to the SVI, the individual policers only affect traffic on the physical ports specified in the secondary interface level of the hierarchical policy map.

A hierarchical policy map has two levels. The first level, the VLAN level, specifies the actions to be taken against a traffic flow on an SVI. The second level, the interface level, specifies the actions to be taken against the traffic on the physical ports that belong to the SVI and are specified in the interface-level policy map.

When configuring policing on an SVI, you can create and configure a hierarchical policy map with these two levels:

- VLAN level—Create this primary level by configuring class maps and classes that specify the port trust state or set a new DSCP or IP precedence value in the packet. The VLAN-level policy map applies only to the VLAN in an SVI and does not support policers.
- Interface level—Create this secondary level by configuring class maps and classes that specify the individual policers on physical ports the belong to the SVI. The interface-level policy map only supports individual policers and does not support aggregate policers. You can configure different interface-level policy maps for each class defined in the VLAN-level policy map.

Figure 38-5 Policing and Marking Flowchart on SVIs

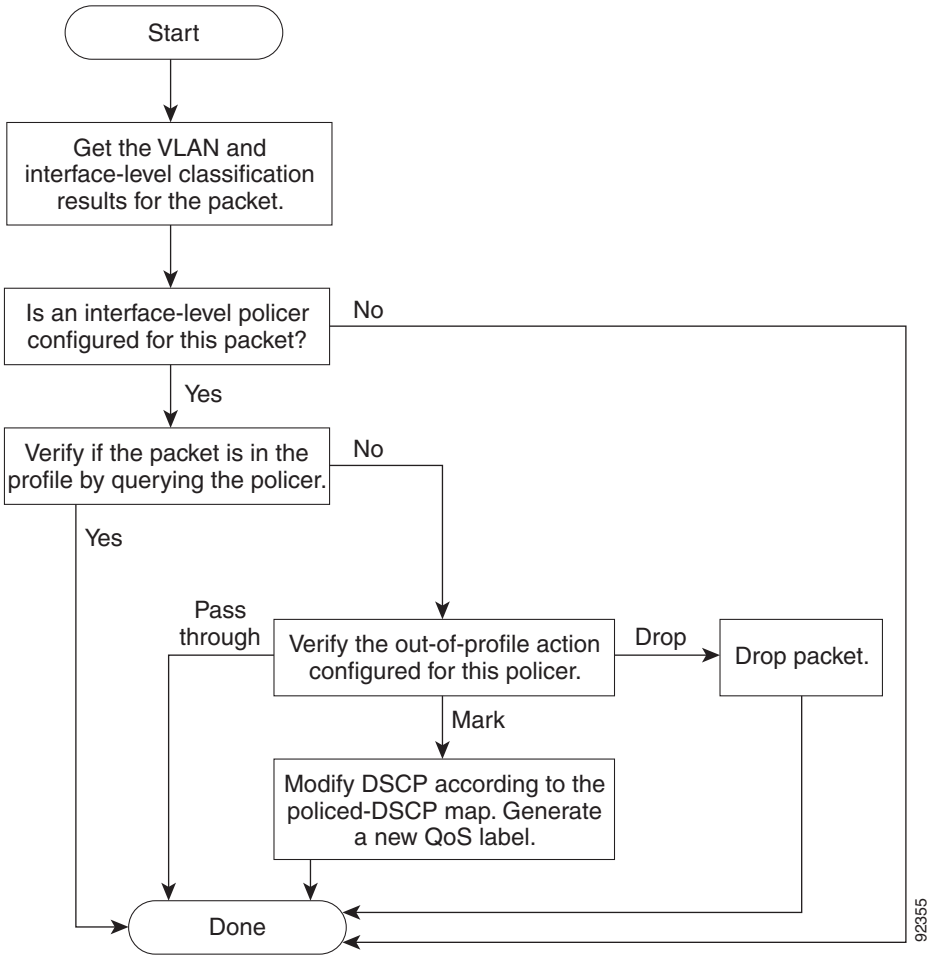

Mapping Tables

During QoS processing, the switch represents the priority of all traffic (including non-IP traffic) with an QoS label based on the DSCP or CoS value from the classification stage:

- During classification, QoS uses configurable mapping tables to derive a corresponding DSCP or CoS value from a received CoS, DSCP, or IP precedence value. These maps include the CoS-to-DSCP map and the IP-precedence-to-DSCP map. You configure these maps by using the **mls qos map cos-dscp** and the **mls qos map ip-prec-dscp** global configuration commands.

  On an ingress port configured in the DSCP-trusted state, if the DSCP values are different between the QoS domains, you can apply the configurable DSCP-to-DSCP-mutation map to the port that is on the boundary between the two QoS domains. You configure this map by using the **mls qos map dscp-mutation** global configuration command.

- During policing, QoS can assign another DSCP value to an IP or a non-IP packet (if the packet is out of profile and the policer specifies a marked-down value). This configurable map is called the policed-DSCP map. You configure this map by using the **mls qos map policed-dscp** global configuration command.

- Before the traffic reaches the scheduling stage, QoS stores the packet in an ingress and an egress queue according to the QoS label. The QoS label is based on the DSCP or the CoS value in the packet and selects the queue through the DSCP input and output queue threshold maps or through the CoS input and output queue threshold maps. In addition to an ingress or an egress queue, the QOS label also identifies the WTD threshold value. You configure these maps by using the **mls qos srr-queue** {**input** | **output**} **dscp-map** and the **mls qos srr-queue** {**input** | **output**} **cos-map** global configuration commands.

The CoS-to-DSCP, DSCP-to-CoS, and the IP-precedence-to-DSCP maps have default values that might or might not be appropriate for your network.

The default DSCP-to-DSCP-mutation map and the default policed-DSCP map are null maps; they map an incoming DSCP value to the same DSCP value. The DSCP-to-DSCP-mutation map is the only map you apply to a specific port. All other maps apply to the entire switch.

For configuration information, see the "Configuring DSCP Maps" section on page 38-47.

For information about the DSCP and CoS input queue threshold maps, see the "Queueing and Scheduling on Ingress Queues" section on page 38-21. For information about the DSCP and CoS output queue threshold maps, see the "Queueing and Scheduling on Egress Queues" section on page 38-22.

Queueing and Scheduling Overview

The switch has queues at specific points to help prevent congestion as shown in Figure 38-6.

Figure 38-6 Ingress and Egress Queue Location

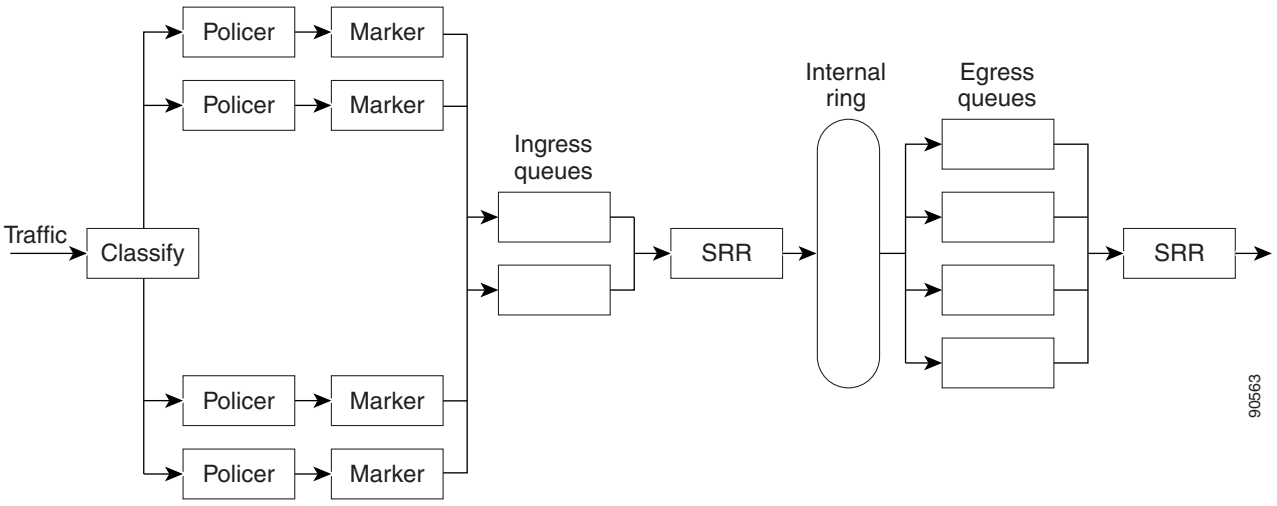

Because the total inbound bandwidth of all ports can exceed the bandwidth of the internal ring, ingress queues are located after the packet is classified, policed, and marked and before packets are forwarded into the switch fabric. Because multiple ingress ports can simultaneously send packets to an egress port and cause congestion, outbound queues are located after the internal ring.

Weighted Tail Drop

Both the ingress and egress queues use an enhanced version of the tail-drop congestion-avoidance mechanism called weighted tail drop (WTD). WTD is implemented on queues to manage the queue lengths and to provide drop precedences for different traffic classifications.

As a frame is enqueued to a particular queue, WTD uses the frame's assigned QoS label to subject it to different thresholds. If the threshold is exceeded for that QoS label (the space available in the destination queue is less than the size of the frame), the switch drops the frame.

Each queue has three threshold values. The QOS label is determines which of the three threshold values is subjected to the frame. Of the three thresholds, two are configurable (explicit) and one is not (implicit).

Figure 38-7 shows an example of WTD operating on a queue whose size is 1000 frames. Three drop percentages are configured: 40 percent (400 frames), 60 percent (600 frames), and 100 percent (1000 frames). These percentages mean that up to 400 frames can be queued at the 40-percent threshold, up to 600 frames at the 60-percent threshold, and up to 1000 frames at the 100-percent threshold.

In this example, CoS values 6 and 7 have a greater importance than the other CoS values, and they are assigned to the 100-percent drop threshold (queue-full state). CoS values 4 and 5 are assigned to the 60-percent threshold, and CoS values 0 to 3 are assigned to the 40-percent threshold.

Suppose the queue is already filled with 600 frames, and a new frame arrives. It contains CoS values 4 and 5 and is subjected to the 60-percent threshold. If this frame is added to the queue, the threshold will be exceeded, so the switch drops it.

Figure 38-7 WTD and Queue Operation

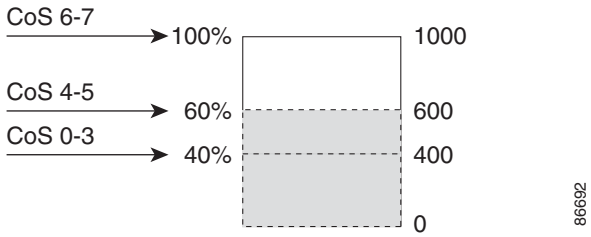

For more information, see the "Mapping DSCP or CoS Values to an Ingress Queue and Setting WTD Thresholds" section on page 38-49, the "Allocating Buffer Space to and Setting WTD Thresholds for an Egress Queue-Set" section on page 38-52, and the "Mapping DSCP or CoS Values to an Egress Queue and to a Threshold ID" section on page 38-53.

SRR Shaping and Sharing

Both the ingress and egress queues are serviced by SRR, which controls the rate at which packets are sent. On the ingress queues, SRR sends packets to the internal ring. On the egress queues, SRR sends packets to the egress port.

You can configure SRR on egress queues for sharing or for shaping. However, for ingress queues, sharing is the default mode, and it is the only mode supported.

In shaped mode, the egress queues are guaranteed a percentage of the bandwidth, and they are rate-limited to that amount. Shaped traffic does not use more than the allocated bandwidth even if the link is idle. Shaping provides a more even flow of traffic over time and reduces the peaks and valleys of bursty traffic. With shaping, the absolute value of each weight is used to compute the bandwidth available for the queues.

In shared mode, the queues share the bandwidth among them according to the configured weights. The bandwidth is guaranteed at this level but not limited to it. For example, if a queue is empty and no longer requires a share of the link, the remaining queues can expand into the unused bandwidth and share it among them. With sharing, the ratio of the weights controls the frequency of dequeuing; the absolute values are meaningless. Shaping and sharing is configured per interface. Each interface can be uniquely configured.

For more information, see the "Allocating Bandwidth Between the Ingress Queues" section on page 38-51, the "Configuring SRR Shaped Weights on Egress Queues" section on page 38-54, and the "Configuring SRR Shared Weights on Egress Queues" section on page 38-55.

Queueing and Scheduling on Ingress Queues

Figure 38-8 shows the queueing and scheduling flowchart for ingress ports.

Figure 38-8 Queueing and Scheduling Flowchart for Ingress Ports

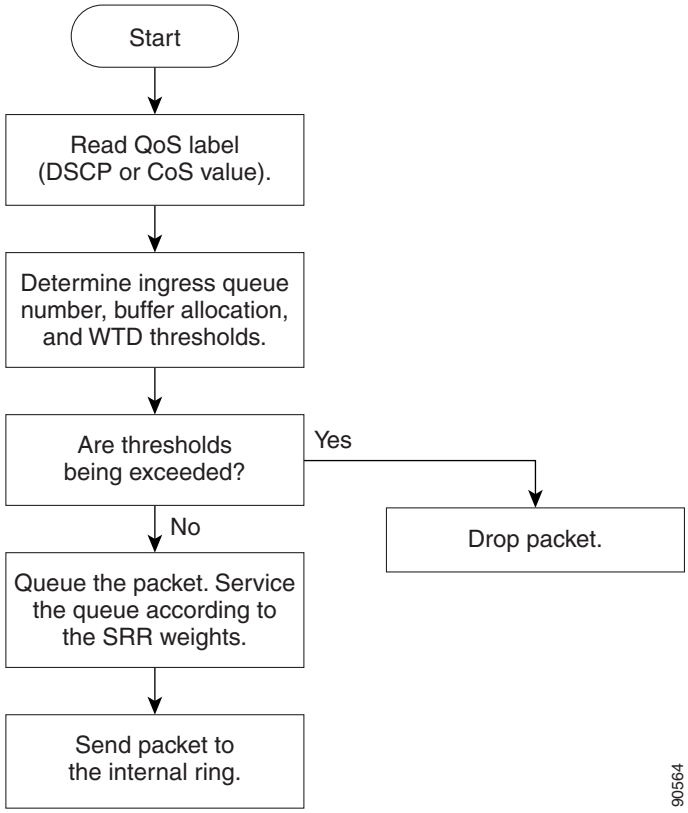

Note SRR services the priority queue for its configured share before servicing the other queue.

The switch supports two configurable ingress queues, which are serviced by SRR in shared mode only. Table 38-10 describes the queues.

Table 38-10 Ingress Queue Types

Queue Type[1]	Function
Normal	User traffic that is considered to be normal priority. You can configure three different thresholds to differentiate among the flows. You can use the **mls qos srr-queue input threshold**, the **mls qos srr-queue input dscp-map**, and the **mls qos srr-queue input cos-map** global configuration commands.
Expedite	High-priority user traffic such as differentiated services (DF) expedited forwarding or voice traffic. You can configure the bandwidth required for this traffic as a percentage of the total traffic by using the **mls qos srr-queue input priority-queue** global configuration command. The expedite queue has guaranteed bandwidth.

1. The switch uses two nonconfigurable queues for traffic that is essential for proper network operation.

You assign each packet that flows through the switch to a queue and to a threshold. Specifically, you map DSCP or CoS values to an ingress queue and map DSCP or CoS values to a threshold ID. You use the **mls qos srr-queue input dscp-map queue** *queue-id* {*dscp1...dscp8* | **threshold** *threshold-id dscp1...dscp8*} or the **mls qos srr-queue input cos-map queue** *queue-id* {*cos1...cos8* | **threshold** *threshold-id cos1...cos8*} global configuration command. You can display the DSCP input queue threshold map and the CoS input queue threshold map by using the **show mls qos maps** privileged EXEC command.

WTD Thresholds

The queues use WTD to support distinct drop percentages for different traffic classes. Each queue has three drop thresholds: two configurable (*explicit*) WTD thresholds and one nonconfigurable (*implicit*) threshold preset to the queue-full state. You assign the two explicit WTD threshold percentages for threshold ID 1 and ID 2 to the ingress queues by using the **mls qos srr-queue input threshold** *queue-id threshold-percentage1 threshold-percentage2* global configuration command. Each threshold value is a percentage of the total number of allocated buffers for the queue. The drop threshold for threshold ID 3 is preset to the queue-full state, and you cannot modify it. For more information about how WTD works, see the "Weighted Tail Drop" section on page 38-19.

Buffer and Bandwidth Allocation

You define the ratio (allocate the amount of space) with which to divide the ingress buffers between the two queues by using the **mls qos srr-queue input buffers** *percentage1 percentage2* global configuration command. The buffer allocation together with the bandwidth allocation control how much data can be buffered and sent before packets are dropped. You allocate bandwidth as a percentage by using the **mls qos srr-queue input bandwidth** *weight1 weight2* global configuration command. The ratio of the weights is the ratio of the frequency in which the SRR scheduler sends packets from each queue.

Priority Queueing

You can configure one ingress queue as the priority queue by using the **mls qos srr-queue input priority-queue** *queue-id* **bandwidth** *weight* global configuration command. The priority queue should be used for traffic (such as voice) that requires guaranteed delivery because this queue is guaranteed part of the bandwidth regardless of the load on the internal ring.

SRR services the priority queue for its configured weight as specified by the **bandwidth** keyword in the **mls qos srr-queue input priority-queue** *queue-id* **bandwidth** *weight* global configuration command. Then, SRR shares the remaining bandwidth with both ingress queues and services them as specified by the weights configured with the **mls qos srr-queue input bandwidth** *weight1 weight2* global configuration command.

You can combine the commands described in this section to prioritize traffic by placing packets with particular DSCPs or CoSs into certain queues, by allocating a large queue size or by servicing the queue more frequently, and by adjusting queue thresholds so that packets with lower priorities are dropped. For configuration information, see the "Configuring Ingress Queue Characteristics" section on page 38-49.

Queueing and Scheduling on Egress Queues

Figure 38-9 shows the queueing and scheduling flowchart for egress ports.

Note If the expedite queue is enabled, SRR services it until it is empty before servicing the other three queues.

Figure 38-9 Queueing and Scheduling Flowchart for Egress Ports

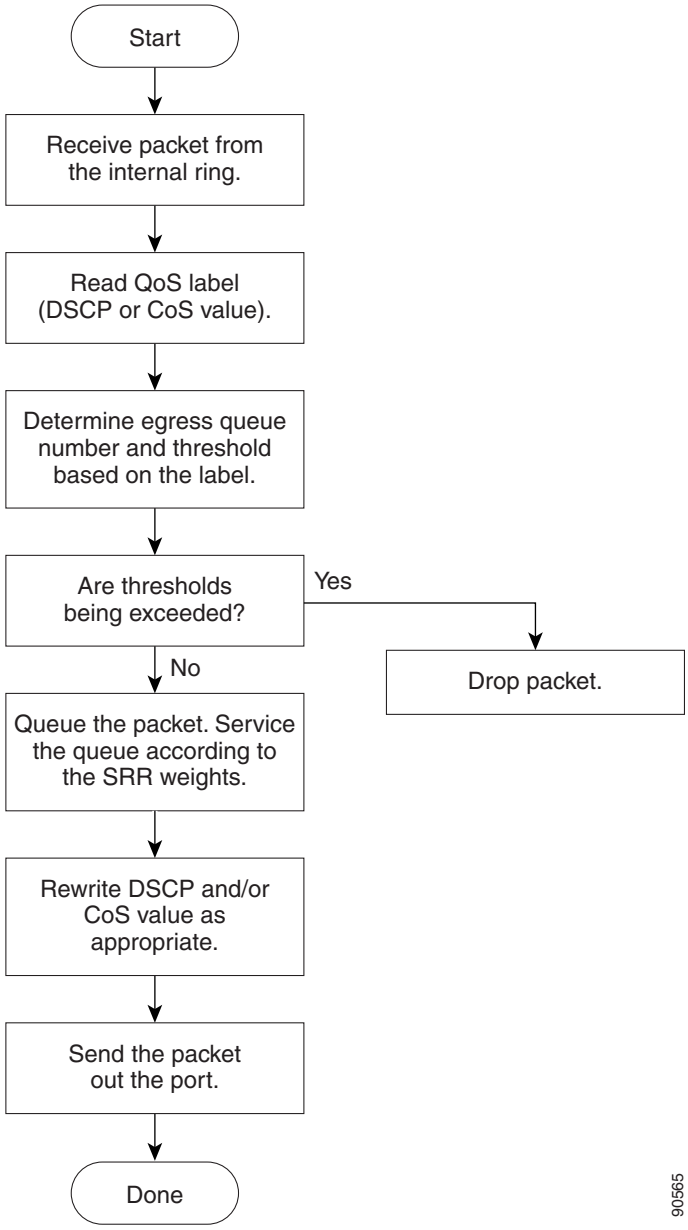

Each port supports four egress queues, one of which (queue 1) can be the egress expedite queue.These queues are configured by a queue-set. All traffic leaving an egress port flows through one of these four queues and is subjected to a threshold based on the QoS label assigned to the packet.

Figure 38-10 shows the egress queue buffer. The buffer space is divided between the common pool and the reserved pool. The switch uses a buffer allocation scheme to reserve a minimum amount of buffers for each egress queue, to prevent any queue or port from consuming all the buffers and depriving other queues, and to control whether to grant buffer space to a requesting queue. The switch detects whether the target queue has not consumed more buffers than its reserved amount (under-limit), whether it has consumed all of its maximum buffers (over limit), and whether the common pool is empty (no free

buffers) or not empty (free buffers). If the queue is not over-limit, the switch can allocate buffer space from the reserved pool or from the common pool (if it is not empty). If there are no free buffers in the common pool or if the queue is over-limit, the switch drops the frame.

Figure 38-10 Egress Queue Buffer Allocation

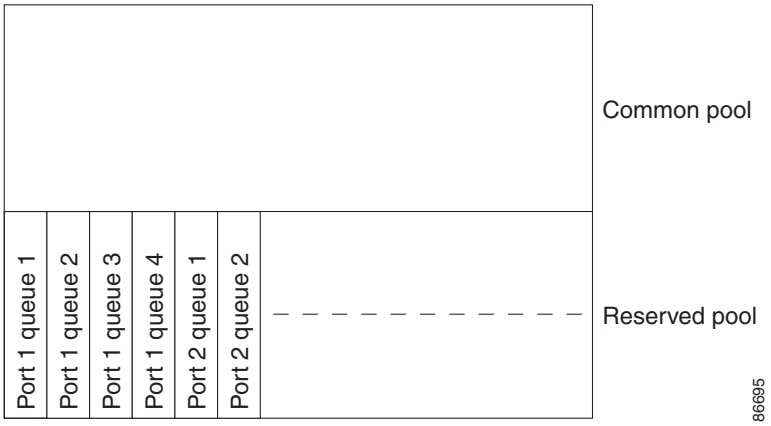

Buffer and Memory Allocation

You guarantee the availability of buffers, set drop thresholds, and configure the maximum memory allocation for a queue-set by using the **mls qos queue-set output** *qset-id* **threshold** *queue-id drop-threshold1 drop-threshold2 reserved-threshold maximum-threshold* global configuration command. Each threshold value is a percentage of the queue's allocated memory, which you specify by using the **mls qos queue-set output** *qset-id* **buffers** *allocation1 ... allocation4* global configuration command. The sum of all the allocated buffers represents the reserved pool, and the remaining buffers are part of the common pool.

Through buffer allocation, you can ensure that high-priority traffic is buffered. For example, if the buffer space is 400, you can allocate 70 percent of it to queue 1 and 10 percent to queues 2 through 4. Queue 1 then has 280 buffers allocated to it, and queues 2 through 4 each have 40 buffers allocated to them.

You can guarantee that the allocated buffers are reserved for a specific queue in a queue-set. For example, if there are 100 buffers for a queue, you can reserve 50 percent (50 buffers). The switch returns the remaining 50 buffers to the common pool. You also can enable a queue in the full condition to obtain more buffers than are reserved for it by setting a maximum threshold. The switch can allocate the needed buffers from the common pool if the common pool is not empty.

WTD Thresholds

You can assign each packet that flows through the switch to a queue and to a threshold. Specifically, you map DSCP or CoS values to an egress queue and map DSCP or CoS values to a threshold ID. You use the **mls qos srr-queue output dscp-map queue** *queue-id* {*dscp1...dscp8* | **threshold** *threshold-id dscp1...dscp8*} or the **mls qos srr-queue output cos-map queue** *queue-id* {*cos1...cos8* | **threshold** *threshold-id cos1...cos8*} global configuration command. You can display the DSCP output queue threshold map and the CoS output queue threshold map by using the **show mls qos maps** privileged EXEC command.

The queues use WTD to support distinct drop percentages for different traffic classes. Each queue has three drop thresholds: two configurable (*explicit*) WTD thresholds and one nonconfigurable (*implicit*) threshold preset to the queue-full state. You assign the two WTD threshold percentages for threshold ID 1 and ID 2. The drop threshold for threshold ID 3 is preset to the queue-full state, and you cannot

modify it. You map a port to queue-set by using the **queue-set qset-id** interface configuration command. Modify the queue-set configuration to change the WTD threshold percentages. For more information about how WTD works, see the "Weighted Tail Drop" section on page 38-19.

Shaped or Shared Mode

SRR services each queue-set in shared or shaped mode. You assign shared or shaped weights to the port by using the **srr-queue bandwidth share** *weight1 weight2 weight3 weight4* or the **srr-queue bandwidth shape** *weight1 weight2 weight3 weight4* interface configuration commands. For an explanation of the differences between shaping and sharing, see the "SRR Shaping and Sharing" section on page 38-20.

The buffer allocation together with the SRR weight ratios control how much data can be buffered and sent before packets are dropped. The weight ratio is the ratio of the frequency in which the SRR scheduler sends packets from each queue.

All four queues participate in the SRR unless the expedite queue is enabled, in which case the first bandwidth weight is ignored and is not used in the ratio calculation. The expedite queue is a priority queue, and it is serviced until empty before the other queues are serviced. You enable the expedite queue by using the **priority-queue out** interface configuration command.

You can combine the commands described in this section to prioritize traffic by placing packets with particular DSCPs or CoSs into certain queues, by allocating a large queue size or by servicing the queue more frequently, and by adjusting queue thresholds so that packets with lower priorities are dropped. For configuration information, see the "Configuring Egress Queue Characteristics" section on page 38-52.

Note The egress queue default settings are suitable for most situations. You should change them only when you have a thorough understanding of the egress queues and if these settings do not meet your QoS solution.

Packet Modification

A packet is classified, policed, and queued to provide QoS. Packet modifications can occur during this process:

- For IP and non-IP packets, classification involves assigning a QoS label to a packet based on the DSCP or CoS of the received packet. However, the packet is not modified at this stage; only an indication of the assigned DSCP or CoS value is carried along. The reason for this is that QoS classification and forwarding lookups occur in parallel, and it is possible that the packet is forwarded with its original DSCP to the CPU where it is again processed through software.
- During policing, IP and non-IP packets can have another DSCP assigned to them (if they are out of profile and the policer specifies a markdown DSCP). Once again, the DSCP in the packet is not modified, but an indication of the marked-down value is carried along. For IP packets, the packet modification occurs at a later stage; for non-IP packets the DSCP is converted to CoS and used for queueing and scheduling decisions.
- Depending on the QoS label assigned to a frame and the mutation chosen, the DSCP and CoS values of the frame are rewritten. If you do not configure the mutation map and if you configure the port to trust the DSCP of the incoming frame, the DSCP value in the frame is not changed, but the CoS is rewritten according to the DSCP-to-CoS map. If you configure the port to trust the CoS of the incoming frame and it is an IP packet, the CoS value in the frame is not changed, but the DSCP might be changed according to the CoS-to-DSCP map.

  The input mutation causes the DSCP to be rewritten depending on the new value of DSCP chosen. The set action in a policy map also causes the DSCP to be rewritten.

Classification Using Port Trust States

Trust State on Ports within the QoS Domain

Packets entering a QoS domain are classified at the edge of the QoS domain. When the packets are classified at the edge, the switch port within the QoS domain can be configured to one of the trusted states because there is no need to classify the packets at every switch within the QoS domain. Figure 38-11 shows a sample network topology.

Figure 38-11 Port Trusted States within the QoS Domain

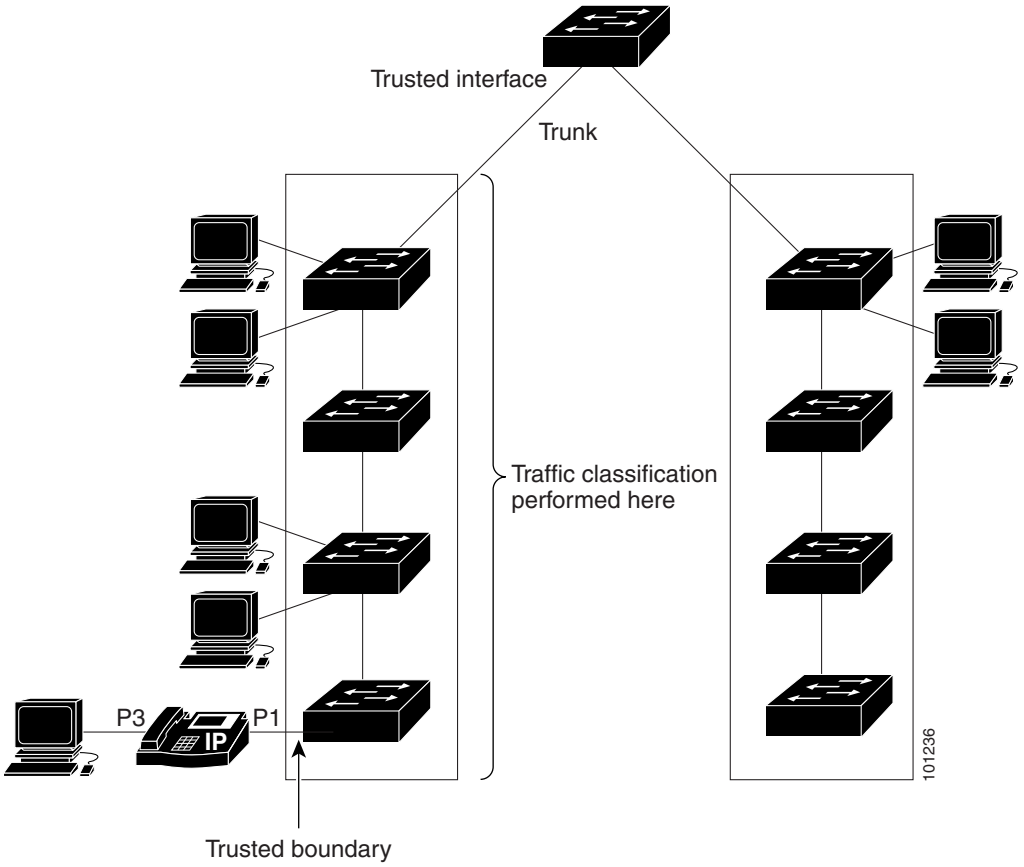

Configuring a Trusted Boundary to Ensure Port Security

In a typical network, you connect a Cisco IP phone to a switch port, as shown in Figure 38-11, and cascade devices that generate data packets from the back of the telephone. The Cisco IP phone guarantees the voice quality through a shared data link by marking the CoS level of the voice packets as high priority (CoS = 5) and by marking the data packets as low priority (CoS = 0). Traffic sent from the telephone to the switch is typically marked with a tag that uses the IEEE 802.1Q header. The header contains the VLAN information and the class of service (CoS) 3-bit field, which is the priority of the packet.

For most Cisco IP phone configurations, the traffic sent from the telephone to the switch should be trusted to ensure that voice traffic is properly prioritized over other types of traffic in the network. By using the **mls qos trust cos** interface configuration command, you configure the switch port to which the telephone is connected to trust the CoS labels of all traffic received on that port. Use the **mls qos trust dscp** interface configuration command to configure a routed port to which the telephone is connected to trust the DSCP labels of all traffic received on that port.

With the trusted setting, you also can use the trusted boundary feature to prevent misuse of a high-priority queue if a user bypasses the telephone and connects the PC directly to the switch. Without trusted boundary, the CoS labels generated by the PC are trusted by the switch (because of the trusted CoS setting). By contrast, trusted boundary uses CDP to detect the presence of a Cisco IP phone (such as the Cisco IP phone 7910, 7935, 7940, and 7960) on a switch port. If the telephone is not detected, the trusted boundary feature disables the trusted setting on the switch port and prevents misuse of a high-priority queue. Note that the trusted boundary feature is not effective if the PC and Cisco IP phone are connected to a hub that is connected to the switch.

In some situations, you can prevent a PC connected to the Cisco IP phone from taking advantage of a high-priority data queue. You can use the **switchport priority extend cos** interface configuration command to configure the telephone through the switch CLI to override the priority of the traffic received from the PC.

DSCP Transparency Mode

The switch supports the DSCP transparency feature. It affects only the DSCP field of a packet at egress. By default, DSCP transparency is disabled. The switch modifies the DSCP field in an incoming packet, and the DSCP field in the outgoing packet is based on the quality of service (QoS) configuration, including the port trust setting, policing and marking, and the DSCP-to-DSCP mutation map.

If DSCP transparency is enabled by using the **no mls qos rewrite ip dscp** command, the switch does not modify the DSCP field in the incoming packet, and the DSCP field in the outgoing packet is the same as that in the incoming packet.

Note Enabling DSCP transparency does not affect the port trust settings on IEEE 802.1Q tunneling ports.

Regardless of the DSCP transparency configuration, the switch modifies the internal DSCP value of the packet, which the switch uses to generate a class of service (CoS) value that represents the priority of the traffic. The switch also uses the internal DSCP value to select an egress queue and threshold.

DSCP Trust State on a Port Bordering Another QoS Domain

If you are administering two separate QoS domains between which you want to implement QoS features for IP traffic, you can configure the switch ports bordering the domains to a DSCP-trusted state as shown in Figure 38-12. Then the receiving port accepts the DSCP-trusted value and avoids the classification stage of QoS. If the two domains use different DSCP values, you can configure the DSCP-to-DSCP-mutation map to translate a set of DSCP values to match the definition in the other domain.

Figure 38-12 DSCP-Trusted State on a Port Bordering Another QoS Domain

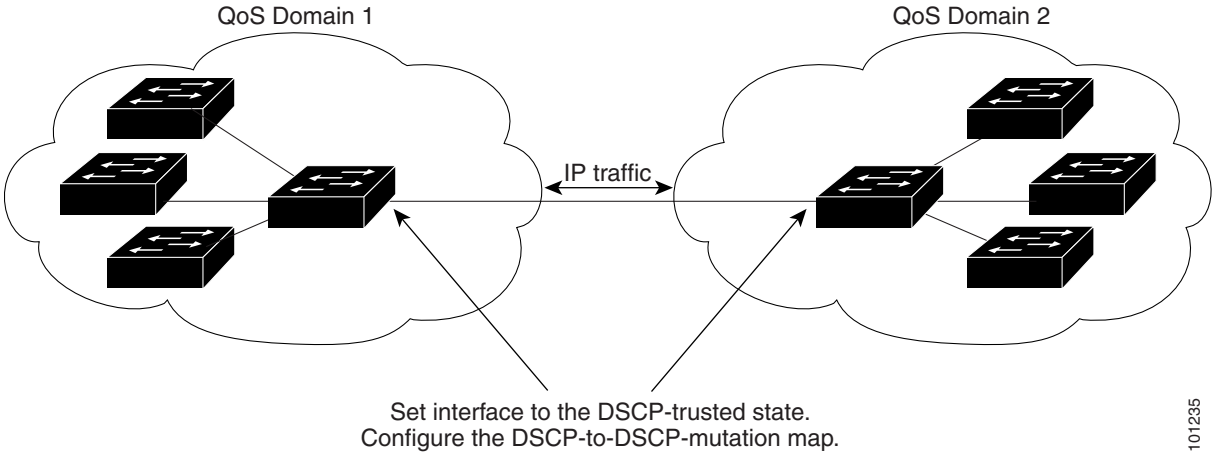

QoS Policies

Classifying, Policing, and Marking Traffic on Physical Ports by Using Policy Maps

You can configure a nonhierarchical policy map on a physical port that specifies which traffic class to act on. Actions can include trusting the CoS, DSCP, or IP precedence values in the traffic class; setting a specific DSCP or IP precedence value in the traffic class; and specifying the traffic bandwidth limitations for each matched traffic class (policer) and the action to take when the traffic is out of profile (marking).

A policy map also has these characteristics:

- A policy map can contain multiple class statements, each with different match criteria and policers.
- A separate policy-map class can exist for each type of traffic received through a port.
- A policy-map trust state and a port trust state are mutually exclusive, and whichever is configured last takes affect.

Follow these guidelines when configuring policy maps on physical ports:

- You can attach only one policy map per ingress port.
- If you configure the IP-precedence-to-DSCP map by using the **mls qos map ip-prec-dscp** *dscp1...dscp8* global configuration command, the settings only affect packets on ingress interfaces that are configured to trust the IP precedence value. In a policy map, if you set the packet IP precedence value to a new value by using the **set ip precedence** *new-precedence* policy-map class configuration command, the egress DSCP value is not affected by the IP-precedence-to-DSCP map. If you want the egress DSCP value to be different than the ingress value, use the **set dscp** *new-dscp* policy-map class configuration command.
- If you enter or have used the **set ip dscp** command, the switch changes this command to **set dscp** in its configuration.
- You can use the **set ip precedence** or the **set precedence** policy-map class configuration command to change the packet IP precedence value. This setting appears as **set ip precedence** in the switch configuration.

- You can configure a separate second-level policy map for each class defined for the port. The second-level policy map specifies the police action to take for each traffic class. For information on configuring a hierarchical policy map, see Classifying, Policing, and Marking Traffic on SVIs by Using Hierarchical Policy Maps, page 38-29.
- A policy map and a port trust state can both run on a physical interface. The policy map is applied before the port trust state.

Classifying, Policing, and Marking Traffic on SVIs by Using Hierarchical Policy Maps

You can configure hierarchical policy maps on SVIs, but not on other types of interfaces. Hierarchical policing combines the VLAN- and interface-level policy maps to create a single policy map.

On an SVI, the VLAN-level policy map specifies which traffic class to act on. Actions can include trusting the CoS, DSCP, or IP precedence values or setting a specific DSCP or IP precedence value in the traffic class. Use the interface-level policy map to specify the physical ports that are affected by individual policers.

Follow these guidelines when configuring hierarchical policy maps:

- Before configuring a hierarchical policy map, you must enable VLAN-based QoS on the physical ports that are to be specified at the interface level of the policy map.
- You can attach only one policy map per ingress port or SVI.
- A policy map can contain multiple class statements, each with different match criteria and actions.
- A separate policy-map class can exist for each type of traffic received on the SVI.
- A policy-map and a port trust state can both run on a physical interface. The policy-map is applied before the port trust state.
- If you configure the IP-precedence-to-DSCP map by using the **mls qos map ip-prec-dscp** *dscp1...dscp8* global configuration command, the settings only affect packets on ingress interfaces that are configured to trust the IP precedence value. In a policy map, if you set the packet IP precedence value to a new value by using the **set ip precedence** *new-precedence* policy-map class configuration command, the egress DSCP value is not affected by the IP-precedence-to-DSCP map. If you want the egress DSCP value to be different than the ingress value, use the **set dscp** *new-dscp* policy-map class configuration command.
- If you enter or have used the **set ip dscp** command, the switch changes this command to **set dscp** in its configuration. If you enter the **set ip dscp** command, this setting appears as **set dscp** in the switch configuration.
- You can use the **set ip precedence** or the **set precedence** policy-map class configuration command to change the packet IP precedence value. This setting appears as **set ip precedence** in the switch configuration.
- If VLAN-based QoS is enabled, the hierarchical policy map supersedes the previously configured port-based policy map.
- The hierarchical policy map is attached to the SVI and affects all traffic belonging to the VLAN. The actions specified in the VLAN-level policy map affect the traffic belonging to the SVI. The police action on the port-level policy map affects the ingress traffic on the affected physical interfaces.
- When configuring a hierarchical policy map on trunk ports, the VLAN ranges must not overlap. If the ranges overlap, the actions specified in the policy map affect the incoming and outgoing traffic on the overlapped VLANs.
- Aggregate policers are not supported in hierarchical policy maps.

- When VLAN-based QoS is enabled, the switch supports VLAN-based features, such as the VLAN map.
- You can configure a hierarchical policy map only on the primary VLAN of a private VLAN.

DSCP Maps

For default DSCP mapping, see "Default Mapping Table Settings" section on page 38-8.

DSCP-to-DSCP-Mutation Map

If two QoS domains have different DSCP definitions, use the DSCP-to-DSCP-mutation map to translate one set of DSCP values to match the definition of another domain. You apply the DSCP-to-DSCP-mutation map to the receiving port (ingress mutation) at the boundary of a QoS administrative domain.

With ingress mutation, the new DSCP value overwrites the one in the packet, and QoS treats the packet with this new value. The switch sends the packet out the port with the new DSCP value.

You can configure multiple DSCP-to-DSCP-mutation maps on an ingress port. The default DSCP-to-DSCP-mutation map is a null map, which maps an incoming DSCP value to the same DSCP value.

Ingress Queue Characteristics

Depending on the complexity of your network and your QoS solution, you might need to perform all of the tasks in the next sections. You will need to make decisions about these characteristics:

- Which packets are assigned (by DSCP or CoS value) to each queue?
- What drop percentage thresholds apply to each queue, and which CoS or DSCP values map to each threshold?
- How much of the available buffer space is allocated between the queues?
- How much of the available bandwidth is allocated between the queues?
- Is there traffic (such as voice) that should be given high priority?

Ingress Priority Queue

You should use the priority queue only for traffic that needs to be expedited (for example, voice traffic, which needs minimum delay and jitter).

The priority queue is guaranteed part of the bandwidth to reduce the delay and jitter under heavy network traffic on an oversubscribed ring (when there is more traffic than the backplane can carry, and the queues are full and dropping frames).

SRR services the priority queue for its configured weight as specified by the **bandwidth** keyword in the **mls qos srr-queue input priority-queue** *queue-id* **bandwidth** *weight* global configuration command. Then, SRR shares the remaining bandwidth with both ingress queues and services them as specified by the weights configured with the **mls qos srr-queue input bandwidth** *weight1 weight2* global configuration command.

Egress Queue Characteristics

Depending on the complexity of your network and your QoS solution, you might need to perform all of the tasks in the next sections. You will need to make decisions about these characteristics:

- Which packets are mapped by DSCP or CoS value to each queue and threshold ID?
- What drop percentage thresholds apply to the queue-set (four egress queues per port), and how much reserved and maximum memory is needed for the traffic type?
- How much of the fixed buffer space is allocated to the queue-set?
- Does the bandwidth of the port need to be rate limited?
- How often should the egress queues be serviced and which technique (shaped, shared, or both) should be used?

Egress Queue Configuration Guidelines

Follow these guidelines when the expedite queue is enabled or the egress queues are serviced based on their SRR weights:

- If the egress expedite queue is enabled, it overrides the SRR shaped and shared weights for queue 1.
- If the egress expedite queue is disabled and the SRR shaped and shared weights are configured, the shaped mode overrides the shared mode for queue 1, and SRR services this queue in shaped mode.
- If the egress expedite queue is disabled and the SRR shaped weights are not configured, SRR services this queue in shared mode.

Allocating Buffer Space to and Setting WTD Thresholds for an Egress Queue-Set

You can guarantee the availability of buffers, set WTD thresholds, and configure the maximum allocation for a queue-set by using the **mls qos queue-set output** *qset-id* **threshold** *queue-id drop-threshold1 drop-threshold2 reserved-threshold maximum-threshold* global configuration commands.

Each threshold value is a percentage of the queues allocated buffers, which you specify by using the **mls qos queue-set output** *qset-id* **buffers** *allocation1 ... allocation4* global configuration command. The queues use WTD to support distinct drop percentages for different traffic classes.

Note The egress queue default settings are suitable for most situations. You should change them only when you have a thorough understanding of the egress queues and if these settings do not meet your QoS solution.

How to Configure Standard QoS

Enabling QoS Globally

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mls qos**	Enables QoS globally.
Step 3	**end**	Returns to privileged EXEC mode.

Enabling VLAN-Based QoS on Physical Ports

By default, VLAN-based QoS is disabled on all physical switch ports. The switch applies QoS, including class maps and policy maps, only on a physical-port basis. You can enable VLAN-based QoS on a switch port.

This procedure is required on physical ports that are specified in the interface level of a hierarchical policy map on an SVI.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the physical port, and enters interface configuration mode.
Step 3	**mls qos vlan-based**	Enables VLAN-based QoS on the port.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring Classification Using Port Trust States

These sections describe how to classify incoming traffic by using port trust states. Depending on your network configuration, you must perform one or more of these tasks or one or more of the tasks in the "Configuring a QoS Policy" section on page 38-36:

- Configuring the Trust State on Ports Within the QoS Domain, page 38-33
- Configuring the CoS Value for an Interface, page 38-33
- Configuring a Trusted Boundary to Ensure Port Security, page 38-34
- Enabling DSCP Transparency Mode, page 38-34
- Configuring the DSCP Trust State on a Port Bordering Another QoS Domain, page 38-35

Configuring the Trust State on Ports Within the QoS Domain

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the port to be trusted, and enters interface configuration mode. Valid interfaces include physical ports.
Step 3	**mls qos trust** [**cos** | **dscp** | **ip-precedence**]	Configures the port trust state. By default, the port is not trusted. If no keyword is specified, the default is **dscp**. The keywords have these meanings: • **cos**—Classifies an ingress packet by using the packet CoS value. For an untagged packet, the port default CoS value is used. The default port CoS value is 0. • **dscp**—Classifies an ingress packet by using the packet DSCP value. For a non-IP packet, the packet CoS value is used if the packet is tagged; for an untagged packet, the default port CoS is used. Internally, the switch maps the CoS value to a DSCP value by using the CoS-to-DSCP map. • **ip-precedence**—Classifies an ingress packet by using the packet IP-precedence value. For a non-IP packet, the packet CoS value is used if the packet is tagged; for an untagged packet, the default port CoS is used. Internally, the switch maps the CoS value to a DSCP value by using the CoS-to-DSCP map.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring the CoS Value for an Interface

QoS assigns the CoS value specified with the **mls qos cos** interface configuration command to untagged frames received on trusted and untrusted ports.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the port to be configured, and enters interface configuration mode. Valid interfaces include physical ports.

	Command	Purpose
Step 3	**mls qos cos** {*default-cos* | **override**}	Configures the default CoS value for the port. • *default-cos*—Specifies a default CoS value to be assigned to a port. If the packet is untagged, the default CoS value becomes the packet CoS value. The CoS range is 0 to 7. The default is 0. • **override**—Overrides the previously configured trust state of the incoming packet and applies the default port CoS value to the port on all incoming packets. By default, CoS override is disabled. Use the **override** keyword when all incoming packets on specified ports deserve higher or lower priority than packets entering from other ports. Even if a port was previously set to trust DSCP, CoS, or IP precedence, this command overrides the previously configured trust state, and all the incoming CoS values are assigned the default CoS value configured with this command. If an incoming packet is tagged, the CoS value of the packet is modified with the default CoS of the port at the ingress port.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring a Trusted Boundary to Ensure Port Security

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**cdp run**	Enables CDP globally. By default, CDP is enabled.
Step 3	**interface** *interface-id*	Specifies the port connected to the Cisco IP phone, and enters interface configuration mode. Valid interfaces include physical ports.
Step 4	**cdp enable**	Enables CDP on the port. By default, CDP is enabled.
Step 5	**mls qos trust cos** or **mls qos trust dscp**	Configures the switch port to trust the CoS value in traffic received from the Cisco IP phone. Configures the routed port to trust the DSCP value in traffic received from the Cisco IP phone. By default, the port is not trusted.
Step 6	**mls qos trust device cisco-phone**	Specifies that the Cisco IP phone is a trusted device. You cannot enable both trusted boundary and auto-QoS (**auto qos voip** interface configuration command) at the same time; they are mutually exclusive.
Step 7	**end**	Returns to privileged EXEC mode.

Enabling DSCP Transparency Mode

To configure the switch to modify the DSCP value based on the trust setting or on an ACL by disabling DSCP transparency, use the **mls qos rewrite ip dscp** global configuration command.

If you disable QoS by using the **no mls qos** global configuration command, the CoS and DSCP values are not changed (the default QoS setting).

If you enter the **no mls qos rewrite ip dscp** global configuration command to enable DSCP transparency and then enter the **mls qos trust** [**cos** | **dscp**] interface configuration command, DSCP transparency is still enabled.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mls qos**	Enables QoS globally.
Step 3	**no mls qos rewrite ip dscp**	Enables DSCP transparency. The switch is configured to not modify the DSCP field of the IP packet.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring the DSCP Trust State on a Port Bordering Another QoS Domain

To ensure a consistent mapping strategy across both QoS domains, you must perform this procedure on the ports in both domains:

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mls qos map dscp-mutation** *dscp-mutation-name in-dscp* **to** *out-dscp*	Modifies the DSCP-to-DSCP-mutation map. The default DSCP-to-DSCP-mutation map is a null map, which maps an incoming DSCP value to the same DSCP value. • *dscp-mutation-name*—Enters the mutation map name. You can create more than one map by specifying a new name. • *in-dscp*—Enters up to eight DSCP values separated by spaces. Then enter the **to** keyword. • *out-dscp*—Enters a single DSCP value. The DSCP range is 0 to 63.
Step 3	**interface** *interface-id*	Specifies the port to be trusted, and enters interface configuration mode. Valid interfaces include physical ports.
Step 4	**mls qos trust dscp**	Configures the ingress port as a DSCP-trusted port. By default, the port is not trusted.
Step 5	**mls qos dscp-mutation** *dscp-mutation-name*	Applies the map to the specified ingress DSCP-trusted port. • *dscp-mutation-name*—Specifies the mutation map name created in Step 2. • You can configure multiple DSCP-to-DSCP-mutation maps on an ingress port.
Step 6	**end**	Returns to privileged EXEC mode.

Configuring a QoS Policy

Configuring a QoS policy typically requires classifying traffic into classes, configuring policies applied to those traffic classes, and attaching policies to ports.

These sections describe how to classify, police, and mark traffic. Depending on your network configuration, you must perform one or more of these tasks:

- Creating IP Standard ACLs, page 38-36
- Creating IP Extended ACLs, page 38-37
- Creating a Layer 2 MAC ACL for Non-IP Traffic, page 38-37
- Creating Class Maps, page 38-38
- Creating Nonhierarchical Policy Maps, page 38-40
- Creating Hierarchical Policy Maps, page 38-42
- Creating Aggregate Policers, page 38-46

Creating IP Standard ACLs

You can classify IP traffic by using IP standard or IP extended ACLs; you can classify non-IP traffic by using Layer 2 MAC ACLs.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*]	Creates an IP standard ACL, repeating the command as many times as necessary. • *access-list-number*—Enters the access list number. The range is 1 to 99 and 1300 to 1999. • **permit**—Permits a certain type of traffic if the conditions are matched. Use the **deny** keyword to deny a certain type of traffic if conditions are matched. • *source*—Enters the network or host from which the packet is being sent. You can use the **any** keyword as an abbreviation for 0.0.0.0 255.255.255.255. • (Optional) *source-wildcard*—Enters the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. **Note** When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.
Step 3	**end**	Returns to privileged EXEC mode.

Creating IP Extended ACLs

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**access-list** *access-list-number* {**deny** | **permit**} *protocol source source-wildcard destination destination-wildcard*	Creates an IP extended ACL, repeating the command as many times as necessary. • *access-list-number*—Enters the access list number. The range is 100 to 199 and 2000 to 2699. • **permit**—Permits a certain type of traffic if the conditions are matched. Use the **deny** keyword to deny a certain type of traffic if conditions are matched. • *protocol*—Enters the name or number of an IP protocol. Use the question mark (?) to see a list of available protocol keywords. • *source*—Enters the network or host from which the packet is being sent. You specify this by using dotted decimal notation, by using the **any** keyword as an abbreviation for *source* 0.0.0.0 *source-wildcard* 255.255.255.255, or by using the **host** keyword for *source* 0.0.0.0. • *source-wildcard*—Enters the wildcard bits by placing ones in the bit positions that you want to ignore. You specify the wildcard by using dotted decimal notation, by using the **any** keyword as an abbreviation for *source* 0.0.0.0 *source-wildcard* 255.255.255.255, or by using the **host** keyword for *source* 0.0.0.0. • *destination*—Enters the network or host to which the packet is being sent. You have the same options for specifying the *destination and destination-wildcard* as those described by *source* and *source-wildcard*. **Note** When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.
Step 3	**end**	Returns to privileged EXEC mode.

Creating a Layer 2 MAC ACL for Non-IP Traffic

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mac access-list extended** *name*	Creates a Layer 2 MAC ACL by specifying the name of the list. After entering this command, the mode changes to extended MAC ACL configuration.

	Command	Purpose
Step 3	{**permit** | **deny**} {**host** *src-MAC-addr mask* | **any** | **host** *dst-MAC-addr* | *dst-MAC-addr mask*} [*type mask*]	Specifies the type of traffic to permit or deny if the conditions are matched, entering the command as many times as necessary. • *src-MAC-addr*—Enters the MAC address of the host from which the packet is being sent. You specify this by using the hexadecimal format (H.H.H), by using the **any** keyword as an abbreviation for *source* 0.0.0, *source-wildcard* ffff.ffff.ffff, or by using the **host** keyword for *source* 0.0.0. • *mask*—Enters the wildcard bits by placing ones in the bit positions that you want to ignore. • *dst-MAC-addr*—Enters the MAC address of the host to which the packet is being sent. You specify this by using the hexadecimal format (H.H.H), by using the **any** keyword as an abbreviation for *source* 0.0.0, *source-wildcard* ffff.ffff.ffff, or by using the **host** keyword for *source* 0.0.0. • (Optional) *type mask*—Specifies the Ethertype number of a packet with Ethernet II or SNAP encapsulation to identify the protocol of the packet. For *type*, the range is from 0 to 65535, typically specified in hexadecimal. For *mask*, enter the *don't care* bits applied to the Ethertype before testing for a match. **Note** When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.
Step 4	**end**	Returns to privileged EXEC mode.

Creating Class Maps

You use the **class-map** global configuration command to name and to isolate a specific traffic flow (or class) from all other traffic. The class map defines the criteria to use to match against a specific traffic flow to further classify it. Match statements can include criteria such as an ACL, IP precedence values, or DSCP values. The match criterion is defined with one match statement entered within the class-map configuration mode.

Note You can also create class-maps during policy map creation by using the **class** policy-map configuration command.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*] or **access-list** *access-list-number* {**deny** | **permit**} *protocol source* [*source-wildcard*] *destination* [*destination-wildcard*] or **mac access-list extended** *name* {**permit** | **deny**} {**host** *src-MAC-addr mask* | **any** | **host** *dst-MAC-addr* | *dst-MAC-addr mask*} [*type mask*]	Creates an IP standard or extended ACL for IP traffic or a Layer 2 MAC ACL for non-IP traffic, repeating the command as many times as necessary. For more information, see the "Creating IP Standard ACLs" section on page 38-36. **Note** When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.
Step 3	**class-map** [**match-all** | **match-any**] *class-map-name*	Creates a class map, and enters class-map configuration mode. By default, no class maps are defined. • (Optional) **match-all**—Performs a logical-AND of all matching statements under this class map. All match criteria in the class map must be matched. • (Optional) **match-any**—Performs a logical-OR of all matching statements under this class map. One or more match criteria must be matched. • *class-map-name*—Specifies the name of the class map. If neither the **match-all** or **match-any** keyword is specified, the default is **match-all**. **Note** Because only one **match** command per class map is supported, the **match-all** and **match-any** keywords function the same.
Step 4	**match** {**access-group** *acl-index-or-name* | **ip dscp** *dscp-list* | **ip precedence** *ip-precedence-list*}	Defines the match criterion to classify traffic. By default, no match criterion is defined. Only one match criterion per class map is supported, and only one ACL per class map is supported. • **access-group** *acl-index-or-name*—Specifies the number or name of the ACL created in Step 2. • **ip dscp** *dscp-list*—Enters a list of up to eight IP DSCP values to match against incoming packets. Separate each value with a space. The range is 0 to 63. • **ip precedence** *ip-precedence-list*—Enters a list of up to eight IP-precedence values to match against incoming packets. Separate each value with a space. The range is 0 to 7.
Step 5	**end**	Returns to privileged EXEC mode.

Creating Nonhierarchical Policy Maps

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**class-map** [**match-all** | **match-any**] *class-map-name*	Creates a class map, and enters class-map configuration mode. By default, no class maps are defined. • (Optional) **match-all**—Performs a logical-AND of all matching statements under this class map. All match criteria in the class map must be matched. • (Optional) **match-any** keyword—Performs a logical-OR of all matching statements under this class map. One or more match criteria must be matched. • *class-map-name*—Specifies the name of the class map. If neither the **match-all** or **match-any** keyword is specified, the default is **match-all**. **Note** Because only one **match** command per class map is supported, the **match-all** and **match-any** keywords function the same.
Step 3	**policy-map** *policy-map-name*	Creates a policy map by entering the policy map name, and enters policy-map configuration mode. By default, no policy maps are defined. The default behavior of a policy map is to set the DSCP to 0 if the packet is an IP packet and to set the CoS to 0 if the packet is tagged. No policing is performed.
Step 4	**class** *class-map-name*	Defines a traffic classification, and enters policy-map class configuration mode. By default, no policy map class-maps are defined. If a traffic class has already been defined by using the **class-map** global configuration command, Specifies its name for *class-map-name* in this command.

	Command	Purpose
Step 5	**trust** [**cos** | **dscp** | **ip-precedence**]	Configures the trust state, which QoS uses to generate a CoS-based or DSCP-based QoS label. **Note** This command is mutually exclusive with the **set** command within the same policy map. If you enter the **trust** command, go to Step 6. By default, the port is not trusted. If no keyword is specified when the command is entered, the default is **dscp**. The keywords have these meanings: • **cos—**QoS derives the DSCP value by using the received or default port CoS value and the CoS-to-DSCP map. • **dscp—**QoS derives the DSCP value by using the DSCP value from the ingress packet. For non-IP packets that are tagged, QoS derives the DSCP value by using the received CoS value; for non-IP packets that are untagged, QoS derives the DSCP value by using the default port CoS value. In either case, the DSCP value is derived from the CoS-to-DSCP map. • **ip-precedence—**QoS derives the DSCP value by using the IP precedence value from the ingress packet and the IP-precedence-to-DSCP map. For non-IP packets that are tagged, QoS derives the DSCP value by using the received CoS value; for non-IP packets that are untagged, QoS derives the DSCP value by using the default port CoS value. In either case, the DSCP value is derived from the CoS-to-DSCP map. For more information, see the "Configuring the CoS-to-DSCP Map" section on page 38-47.
Step 6	**set** {**dscp** *new-dscp* | **ip precedence** *new-precedence*}	Classifies IP traffic by setting a new value in the packet. • **dscp** *new-dscp*—Enters a new DSCP value to be assigned to the classified traffic. The range is 0 to 63. • **ip precedence** *new-precedence*—Enters a new IP-precedence value to be assigned to the classified traffic. The range is 0 to 7.
Step 7	**police** *rate-bps burst-byte* [**exceed-action** {**drop** | **policed-dscp-transmit**}]	Defines a policer for the classified traffic. By default, no policer is defined. For information on the number of policers supported, see the "Standard QoS Configuration Guidelines" section on page 38-5. • *rate-bps*—Specifies average traffic rate in bits per second (b/s). The range is 8000 to 1000000000. • *burst-byte*—Specifies the normal burst size in bytes. The range is 8000 to 1000000. • (Optional) Specifies the action to take when the rates are exceeded. Use the **exceed-action drop** keywords to drop the packet. Use the **exceed-action policed-dscp-transmit** keywords to mark down the DSCP value (by using the policed-DSCP map) and to send the packet. For more information, see the "Configuring the Policed-DSCP Map" section on page 38-48.

	Command	Purpose
Step 8	**end**	Returns to global configuration mode.
Step 9	**interface** *interface-id*	Specifies the port to attach to the policy map, and enters interface configuration mode. Valid interfaces include physical ports.
Step 10	**service-policy input** *policy-map-name*	Specifies the policy-map name, and applies it to an ingress port. Only one policy map per ingress port is supported.
Step 11	**end**	Returns to privileged EXEC mode.

Creating Hierarchical Policy Maps

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**class-map** [**match-all** | **match-any**] *class-map-name*	Creates a VLAN-level class map, and enter class-map configuration mode. By default, no class maps are defined. • (Optional) **match-all**—Performs a logical-AND of all matching statements under this class map. All match criteria in the class map must be matched. • (Optional) **match-any**—Performs a logical-OR of all matching statements under this class map. One or more match criteria must be matched. • *class-map-name*—Specifies the name of the class map. If neither the **match-all** or **match-any** keyword is specified, the default is **match-all**. **Note** Because only one **match** command per class map is supported, the **match-all** and **match-any** keywords function the same.
Step 3	**match** {**access-group** *acl-index-or-name* | **ip dscp** *dscp-list* | **ip precedence** *ip-precedence-list*}	Defines the match criterion to classify traffic. By default, no match criterion is defined. Only one match criterion per class map is supported, and only one ACL per class map is supported. • **access-group** *acl-index-or-name*—Specifies the number or name of the ACL. • **ip dscp** *dscp-list*—Enters a list of up to eight IP DSCP values to match against incoming packets. Separate each value with a space. The range is 0 to 63. • **ip precedence** *ip-precedence-list*—Enters a list of up to eight IP-precedence values to match against incoming packets. Separate each value with a space. The range is 0 to 7.
Step 4	**end**	Returns to global configuration mode.

	Command	Purpose
Step 5	**class-map** [**match-all** | **match-any**] *class-map-name*	Creates an interface-level class map, and enters class-map configuration mode. By default, no class maps are defined. • (Optional) **match-all**—Performs a logical-AND of all matching statements under this class map. All match criteria in the class map must be matched. • (Optional) **match-any**—Performs a logical-OR of all matching statements under this class map. One or more match criteria must be matched. • *class-map-name*—Specifies the name of the class map. If neither the **match-all** or **match-any** keyword is specified, the default is **match-all**. **Note** Because only one **match** command per class map is supported, the **match-all** and **match-any** keywords function the same.
Step 6	**match input-interface** *interface-id-list*	Specifies the physical ports on which the interface-level class map acts. You can specify up to six ports as follows: • A single port (counts as one entry) • A list of ports separated by a space (each port counts as an entry) • A range of ports separated by a hyphen (counts as two entries) This command can only be used in the child-level policy map and must be the only match condition in the child-level policy map.
Step 7	**end**	Returns to global configuration mode.
Step 8	**policy-map** *policy-map-name*	Creates an interface-level policy map by entering the policy-map name, and enter policy-map configuration mode. By default, no policy maps are defined, and no policing is performed.
Step 9	**class-map** *class-map-name*	Defines an interface-level traffic classification, and enters policy-map configuration mode. By default, no policy-map class-maps are defined. If a traffic class has already been defined by using the **class-map** global configuration command, specify its name for *class-map-name* in this command.

	Command	Purpose
Step 10	**police** *rate-bps burst-byte* [**exceed-action** {**drop** | **policed-dscp-transmit**}]	Defines an individual policer for the classified traffic. By default, no policer is defined. For information on the number of policers supported, see the "Standard QoS Configuration Guidelines" section on page 38-5. • *rate-bps*—Specifies average traffic rate in bits per second (b/s). The range is 8000 to 1000000000. • *burst-byte*—Specifies the normal burst size in bytes. The range is 8000 to 1000000. • (Optional)—Specifies the action to take when the rates are exceeded. Use the **exceed-action drop** keywords to drop the packet. Use the **exceed-action policed-dscp-transmit** keywords to mark down the DSCP value (by using the policed-DSCP map) and to send the packet. For more information, see the "Configuring the Policed-DSCP Map" section on page 38-48**.**
Step 11	**exit**	Returns to global configuration mode.
Step 12	**policy-map** *policy-map-name*	Creates a VLAN-level policy map by entering the policy-map name, and enters policy-map configuration mode. By default, no policy maps are defined. The default behavior of a policy map is to set the DSCP to 0 if the packet is an IP packet and to set the CoS to 0 if the packet is tagged. No policing is performed.
Step 13	**class** *class-map-name*	Defines a VLAN-level traffic classification, and enter policy-map class configuration mode. By default, no policy-map class-maps are defined. If a traffic class has already been defined by using the **class-map** global configuration command, specify its name for *class-map-name* in this command.

	Command	Purpose
Step 14	**trust** [**cos** | **dscp** | **ip-precedence**]	Configures the trust state, which QoS uses to generate a CoS-based or DSCP-based QoS label. **Note** This command is mutually exclusive with the **set** command within the same policy map. If you enter the **trust** command, omit Step 18. By default, the port is not trusted. If no keyword is specified when the command is entered, the default is **dscp**. The keywords have these meanings: • **cos**—QoS derives the DSCP value by using the received or default port CoS value and the CoS-to-DSCP map. • **dscp**—QoS derives the DSCP value by using the DSCP value from the ingress packet. For non-IP packets that are tagged, QoS derives the DSCP value by using the received CoS value; for non-IP packets that are untagged, QoS derives the DSCP value by using the default port CoS value. In either case, the DSCP value is derived from the CoS-to-DSCP map. • **ip-precedence**—QoS derives the DSCP value by using the IP precedence value from the ingress packet and the IP-precedence-to-DSCP map. For non-IP packets that are tagged, QoS derives the DSCP value by using the received CoS value; for non-IP packets that are untagged, QoS derives the DSCP value by using the default port CoS value. In either case, the DSCP value is derived from the CoS-to-DSCP map. For more information, see the "Configuring the CoS-to-DSCP Map" section on page 38-47.
Step 15	**set** {**dscp** *new-dscp* | **ip precedence** *new-precedence*}	Classifies IP traffic by setting a new value in the packet. • **dscp** *new-dscp*—Enters a new DSCP value to be assigned to the classified traffic. The range is 0 to 63. • **ip precedence** *new-precedence*—Enters a new IP-precedence value to be assigned to the classified traffic. The range is 0 to 7.
Step 16	**service-policy** *policy-map-name*	Specifies the interface-level policy-map name (from Step 10) and associate it with the VLAN-level policy map. If the VLAN-level policy map specifies more than one class, each class can have a different **service-policy** *policy-map-name* command.
Step 17	**end**	Returns to global configuration mode.
Step 18	**interface** *interface-id*	Specifies the SVI to which to attach the hierarchical policy map, and enter interface configuration mode.

	Command	Purpose
Step 19	**service-policy input** *policy-map-name*	Specifies the VLAN-level policy-map name, and applies it to the SVI. Repeat the previous step and this command to apply the policy map to other SVIs. If the hierarchical VLAN-level policy map has more than one interface-level policy map, all class maps must be configured to the same VLAN-level policy map specified in the **service-policy** *policy-map-name* command.
Step 20	**end**	Returns to privileged EXEC mode.

Creating Aggregate Policers

By using an aggregate policer, you can create a policer that is shared by multiple traffic classes within the same policy map. However, you cannot use the aggregate policer across different policy maps or ports.

You can configure aggregate policers only in nonhierarchical policy maps on physical ports.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mls qos aggregate-policer** *aggregate-policer-name rate-bps burst-byte* **exceed-action** {**drop** | **policed-dscp-transmit**}	Defines the policer parameters that can be applied to multiple traffic classes within the same policy map. By default, no aggregate policer is defined. For information on the number of policers supported, see the "Standard QoS Configuration Guidelines" section on page 38-5. • *aggregate-policer-name*—Specifies the name of the aggregate policer. • *ate-bps*—Specifies average traffic rate in bits per second (b/s). The range is 8000 to 1000000000. • *burst-byte*—Specifies the normal burst size in bytes. The range is 8000 to 1000000. • Specifies the action to take when the rates are exceeded. Use the **exceed-action drop** keywords to drop the packet. Use the **exceed-action policed-dscp-transmit** keywords to mark down the DSCP value (by using the policed-DSCP map) and to send the packet.
Step 3	**class-map** [**match-all** | **match-any**] *class-map-name*	Creates a class map to classify traffic as necessary.
Step 4	**policy-map** *policy-map-name*	Creates a policy map by entering the policy map name, and enters policy-map configuration mode.
Step 5	**class** *class-map-name*	Defines a traffic classification, and enters policy-map class configuration mode.
Step 6	**police aggregate** *aggregate-policer-name*	Applies an aggregate policer to multiple classes in the same policy map. • *aggregate-policer-name*—Enters the name specified in Step 2.

	Command	Purpose
Step 7	**exit**	Returns to global configuration mode.
Step 8	**interface** *interface-id*	Specifies the port to attach to the policy map, and enters interface configuration mode. Valid interfaces include physical ports.
Step 9	**service-policy input** *policy-map-name*	Specifies the policy-map name, and applies it to an ingress port. Only one policy map per ingress port is supported.
Step 10	**end**	Returns to privileged EXEC mode.

Configuring DSCP Maps

These sections contain this configuration information:

- Configuring the CoS-to-DSCP Map, page 38-47 (optional)
- Configuring the IP-Precedence-to-DSCP Map, page 38-48 (optional)
- Configuring the Policed-DSCP Map, page 38-48 (optional, unless the null settings in the map are not appropriate)
- Configuring the DSCP-to-CoS Map, page 38-48 (optional)
- Configuring the DSCP-to-DSCP-Mutation Map, page 38-49 (optional, unless the null settings in the map are not appropriate)

For the default DSCP maps, see "Default Mapping Table Settings" section on page 38-8.

All the maps, except the DSCP-to-DSCP-mutation map, are globally defined and are applied to all ports.

Configuring the CoS-to-DSCP Map

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mls qos map cos-dscp** *dscp1...dscp8*	Modifies the CoS-to-DSCP map. For *dscp1...dscp8*, enter eight DSCP values that correspond to CoS values 0 to 7. Separate each DSCP value with a space. The DSCP range is 0 to 63.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring the IP-Precedence-to-DSCP Map

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mls qos map ip-prec-dscp** *dscp1...dscp8*	Modifies the IP-precedence-to-DSCP map. • *dscp1...dscp8*—Enters eight DSCP values that correspond to the IP precedence values 0 to 7. Separate each DSCP value with a space. The DSCP range is 0 to 63.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring the Policed-DSCP Map

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mls qos map policed-dscp** *dscp-list* **to** *mark-down-dscp*	Modifies the policed-DSCP map. • *dscp-list*—Enters up to eight DSCP values separated by spaces. Then enter the **to** keyword. • *mark-down-dscp*—Enters the corresponding policed (marked down) DSCP value.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring the DSCP-to-CoS Map

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mls qos map dscp-cos** *dscp-list* **to** *cos*	Modifies the DSCP-to-CoS map. • *dscp-list*—Enters up to eight DSCP values separated by spaces and then enter the **to** keyword. • *cos*—Enters the CoS value to which the DSCP values correspond. The DSCP range is 0 to 63; the CoS range is 0 to 7.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring the DSCP-to-DSCP-Mutation Map

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mls qos map dscp-mutation** *dscp-mutation-name in-dscp* **to** *out-dscp*	Modifies the DSCP-to-DSCP-mutation map. • *dscp-mutation-name*—Enters the mutation map name. You can create more than one map by specifying a new name. • *in-dscp*—Enters up to eight DSCP values separated by spaces. Then enter the **to** keyword. • *out-dscp*—Enters a single DSCP value. The DSCP range is 0 to 63.
Step 3	**interface** *interface-id*	Specifies the port to which to attach the map, and enters interface configuration mode. Valid interfaces include physical ports.
Step 4	**mls qos trust dscp**	Configures the ingress port as a DSCP-trusted port. By default, the port is not trusted.
Step 5	**mls qos dscp-mutation** *dscp-mutation-name*	Applies the map to the specified ingress DSCP-trusted port. • *dscp-mutation-name*—Enters the mutation map name specified in Step 2.
Step 6	**end**	Returns to privileged EXEC mode.

Configuring Ingress Queue Characteristics

These sections contain this configuration information:

- Mapping DSCP or CoS Values to an Ingress Queue and Setting WTD Thresholds, page 38-49 (optional)
- Allocating Buffer Space Between the Ingress Queues, page 38-50 (optional)
- Allocating Bandwidth Between the Ingress Queues, page 38-51 (optional)
- Configuring the Ingress Priority Queue, page 38-51 (optional)

Mapping DSCP or CoS Values to an Ingress Queue and Setting WTD Thresholds

You can prioritize traffic by placing packets with particular DSCPs or CoSs into certain queues and adjusting the queue thresholds so that packets with lower priorities are dropped.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mls qos srr-queue input dscp-map queue** *queue-id* **threshold** *threshold-id dscp1...dscp8* or **mls qos srr-queue input cos-map queue** *queue-id* **threshold** *threshold-id cos1...cos8*	Maps DSCP or CoS values to an ingress queue and to a threshold ID. By default, DSCP values 0–39 and 48–63 are mapped to queue 1 and threshold 1. DSCP values 40–47 are mapped to queue 2 and threshold 1. By default, CoS values 0–4, 6, and 7 are mapped to queue 1 and threshold 1. CoS value 5 is mapped to queue 2 and threshold 1. • *queue-id*—The range is 1 to 2. • *threshold-id*—The range is 1 to 3. The drop-threshold percentage for threshold 3 is predefined. It is set to the queue-full state. • *dscp1...dscp8*—Enter up to eight values, and separate each value with a space. The range is 0 to 63. • *cos1...cos8*—Enter up to eight values, and separate each value with a space. The range is 0 to 7.
Step 3	**mls qos srr-queue input threshold** *queue-id threshold-percentage1 threshold-percentage2*	Assigns the two WTD threshold percentages for (threshold 1 and 2) to an ingress queue. The default, both thresholds are set to 100 percent. • *queue-id*—The range is 1 to 2. • *threshold-percentage1 threshold-percentage2*—The range is 1 to 100. Separate each value with a space. Each threshold value is a percentage of the total number of queue descriptors allocated for the queue.
Step 4	**end**	Returns to privileged EXEC mode.

Allocating Buffer Space Between the Ingress Queues

You define the ratio (allocate the amount of space) with which to divide the ingress buffers between the two queues. The buffer and the bandwidth allocation control how much data can be buffered before packets are dropped.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mls qos srr-queue input buffers** *percentage1 percentage2*	Allocates the buffers between the ingress queues By default 90 percent of the buffers are allocated to queue 1, and 10 percent of the buffers are allocated to queue 2. *percentage1 percentage2*—The range is 0 to 100. Separate each value with a space. You should allocate the buffers so that the queues can handle any incoming bursty traffic.
Step 3	**end**	Returns to privileged EXEC mode.

Allocating Bandwidth Between the Ingress Queues

You need to specify how much of the available bandwidth is allocated between the ingress queues. The ratio of the weights is the ratio of the frequency in which the SRR scheduler sends packets from each queue. The bandwidth and the buffer allocation control how much data can be buffered before packets are dropped. On ingress queues, SRR operates only in shared mode.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mls qos srr-queue input bandwidth** *weight1 weight2*	Assigns shared round robin weights to the ingress queues. The default setting for *weight1* and *weight2* is 4 (1/2 of the bandwidth is equally shared between the two queues). *weight1* and *weight2*—the range is 1 to 100. Separate each value with a space. SRR services the priority queue for its configured weight as specified by the **bandwidth** keyword in the **mls qos srr-queue input priority-queue** *queue-id* **bandwidth** *weight* global configuration command. Then, SRR shares the remaining bandwidth with both ingress queues and services them as specified by the weights configured with the **mls qos srr-queue input bandwidth** *weight1 weight2* global configuration command. For more information, see the "Configuring the Ingress Priority Queue" section on page 38-51.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring the Ingress Priority Queue

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mls qos srr-queue input priority-queue** *queue-id* **bandwidth** *weight*	Assigns a queue as the priority queue and guarantees bandwidth on the internal ring if the ring is congested. By default, the priority queue is queue 2, and 10 percent of the bandwidth is allocated to it. • *queue-id*—the range is 1 to 2. • **bandwidth** *weight*—Assigns the bandwidth percentage of the internal ring. The range is 0 to 40. The amount of bandwidth that can be guaranteed is restricted because a large value affects the entire ring and can degrade performance.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring Egress Queue Characteristics

These sections contain this configuration information:

- Allocating Buffer Space to and Setting WTD Thresholds for an Egress Queue-Set, page 38-52
- Allocating Buffer Space to and Setting WTD Thresholds for an Egress Queue-Set, page 38-52 (optional)
- Mapping DSCP or CoS Values to an Egress Queue and to a Threshold ID, page 38-53 (optional)
- Configuring SRR Shaped Weights on Egress Queues, page 38-54 (optional)
- Configuring SRR Shared Weights on Egress Queues, page 38-55 (optional)
- Configuring the Egress Expedite Queue, page 38-56 (optional)
- Limiting the Bandwidth on an Egress Interface, page 38-56 (optional)

Allocating Buffer Space to and Setting WTD Thresholds for an Egress Queue-Set

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mls qos queue-set output** *qset-id* **buffers** *allocation1 ... allocation4*	Allocates buffers to a queue-set. By default, all allocation values are equally mapped among the four queues (25, 25, 25, 25). Each queue has 1/4 of the buffer space. • *qset-id*—Enters the ID of the queue-set. The range is 1 to 2. Each port belongs to a queue-set, which defines all the characteristics of the four egress queues per port. • *allocation1 ... allocation4*—Specifies four percentages, one for each queue in the queue-set. For *allocation1*, *allocation3*, and *allocation4*, the range is 0 to 99. For *allocation2*, the range is 1 to 100 (including the CPU buffer). Allocates buffers according to the importance of the traffic; for example, give a large percentage of the buffer to the queue with the highest-priority traffic.

	Command	Purpose
Step 3	**mls qos queue-set output** *qset-id* **threshold** *queue-id drop-threshold1 drop-threshold2 reserved-threshold maximum-threshold*	Configures the WTD thresholds, guarantees the availability of buffers, and configures the maximum memory allocation for the queue-set (four egress queues per port). By default, the WTD thresholds for queues 1, 3, and 4 are set to 100 percent. The thresholds for queue 2 are set to 200 percent. The reserved thresholds for queues 1, 2, 3, and 4 are set to 50 percent. The maximum thresholds for all queues are set to 400 percent. • *qset-id*—Enters the ID of the queue-set specified in Step 2. The range is 1 to 2. • *queue-id*—Enters the specific queue in the queue-set on which the command is performed. The range is 1 to 4. • *drop-threshold1 drop-threshold2*—Specifies the two WTD thresholds expressed as a percentage of the queue's allocated memory. The range is 1 to 3200 percent. • *reserved-threshold*—Enters the amount of memory to be guaranteed (reserved) for the queue expressed as a percentage of the allocated memory. The range is 1 to 100 percent. • *maximum-threshold*—Enables a queue in the full condition to obtain more buffers than are reserved for it. This is the maximum memory the queue can have before the packets are dropped if the common pool is not empty. The range is 1 to 3200 percent.
Step 4	**interface** *interface-id*	Specifies the port of the outbound traffic, and enters interface configuration mode.
Step 5	**queue-set** *qset-id*	Maps the port to a queue-set. • *qset-id*—Enters the ID of the queue-set specified in Step 2. The range is 1 to 2. The default is 1.
Step 6	**end**	Returns to privileged EXEC mode.

Mapping DSCP or CoS Values to an Egress Queue and to a Threshold ID

You can prioritize traffic by placing packets with particular DSCPs or costs of service into certain queues and adjusting the queue thresholds so that packets with lower priorities are dropped.

Note The egress queue default settings are suitable for most situations. You should change them only when you have a thorough understanding of the egress queues and if these settings do not meet your QoS solution.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mls qos srr-queue output dscp-map queue** *queue-id* **threshold** *threshold-id dscp1...dscp8* or **mls qos srr-queue output cos-map queue** *queue-id* **threshold** *threshold-id cos1...cos8*	Maps DSCP or CoS values to an egress queue and to a threshold ID. By default, DSCP values 0–15 are mapped to queue 2 and threshold 1. DSCP values 16–31 are mapped to queue 3 and threshold 1. DSCP values 32–39 and 48–63 are mapped to queue 4 and threshold 1. DSCP values 40–47 are mapped to queue 1 and threshold 1. By default, CoS values 0 and 1 are mapped to queue 2 and threshold 1. CoS values 2 and 3 are mapped to queue 3 and threshold 1. CoS values 4, 6, and 7 are mapped to queue 4 and threshold 1. CoS value 5 is mapped to queue 1 and threshold 1. • *queue-id*—Specifies the range 1 to 4. • *threshold-id*—Specifies the range 1 to 3. The drop-threshold percentage for threshold 3 is predefined. It is set to the queue-full state. • *dscp1...dscp8*—Enters up to eight values, and separates each value with a space. The range is 0 to 63. • *cos1...cos8*—Enters up to eight values, and separates each value with a space. The range is 0 to 7.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring SRR Shaped Weights on Egress Queues

You can specify how much of the available bandwidth is allocated to each queue. The ratio of the weights is the ratio of frequency in which the SRR scheduler sends packets from each queue.

You can configure the egress queues for shaped or shared weights, or both. Use shaping to smooth bursty traffic or to provide a smoother output over time. For information about shaped weights, see the "SRR Shaping and Sharing" section on page 38-20. For information about shared weights, see the "Configuring SRR Shared Weights on Egress Queues" section on page 38-55.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the port of the outbound traffic, and enters interface configuration mode.

	Command	Purpose
Step 3	**srr-queue bandwidth shape** *weight1 weight2 weight3 weight4*	Assigns SRR weights to the egress queues. By default, weight1 is set to 25; weight2, weight3, and weight4 are set to 0, and these queues are in shared mode. *weight1 weight2 weight3 weight4*—Enters the weights to control the percentage of the port that is shaped. The inverse ratio (1/*weight*) controls the shaping bandwidth for this queue. Separate each value with a space. The range is 0 to 65535. If you configure a weight of 0, the corresponding queue operates in shared mode. The weight specified with the **srr-queue bandwidth shape** command is ignored, and the weights specified with the **srr-queue bandwidth share** interface configuration command for a queue come into effect. When configuring queues in the same queue-set for both shaping and sharing, make sure that you configure the lowest number queue for shaping. The shaped mode overrides the shared mode.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring SRR Shared Weights on Egress Queues

In shared mode, the queues share the bandwidth among them according to the configured weights. The bandwidth is guaranteed at this level but not limited to it. For example, if a queue empties and does not require a share of the link, the remaining queues can expand into the unused bandwidth and share it among them. With sharing, the ratio of the weights controls the frequency of dequeuing; the absolute values are meaningless.

Note The egress queue default settings are suitable for most situations. You should change them only when you have a thorough understanding of the egress queues and if these settings do not meet your QoS solution.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the port of the outbound traffic, and enters interface configuration mode.
Step 3	**srr-queue bandwidth share** *weight1 weight2 weight3 weight4*	Assigns SRR weights to the egress queues. By default, all four weights are 25 (1/4 of the bandwidth is allocated to each queue). • *weight1 weight2 weight3 weight4*—Enters the weights to control the ratio of the frequency in which the SRR scheduler sends packets. Separate each value with a space. The range is 1 to 255.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring the Egress Expedite Queue

You can ensure that certain packets have priority over all others by queuing them in the egress expedite queue. SRR services this queue until it is empty before servicing the other queues.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**mls qos**	Enables QoS on a switch.
Step 3	**interface** *interface-id*	Specifies the egress port, and enters interface configuration mode.
Step 4	**priority-queue out**	Enables the egress expedite queue, which is disabled by default. When you configure this command, the SRR weight and queue size ratios are affected because there is one less queue participating in SRR. This means that *weight1* in the **srr-queue bandwidth shape** or the **srr-queue bandwidth share** command is ignored (not used in the ratio calculation).
Step 5	**end**	Returns to privileged EXEC mode.

Limiting the Bandwidth on an Egress Interface

You can limit the bandwidth on an egress port. For example, if a customer pays only for a small percentage of a high-speed link, you can limit the bandwidth to that amount.

Note The egress queue default settings are suitable for most situations. You should change them only when you have a thorough understanding of the egress queues and if these settings do not meet your QoS solution.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the port to be rate limited, and enters interface configuration mode.
Step 3	**srr-queue bandwidth limit** *weight1*	Specifies the percentage of the port speed to which the port should be limited. The range is 10 to 90. By default, the port is not rate limited and is set to 100 percent.
Step 4	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining Standard QoS

Command	Purpose
show access-lists	Verifies your entries.
show class-map [*class-map-name*]	Displays QoS class maps, which define the match criteria to classify traffic.

Command	Purpose
show mls qos	Displays global QoS configuration information.
show mls qos aggregate-policer [*aggregate-policer-name*]	Displays the aggregate policer configuration.
show mls qos input-queue	Displays QoS settings for the ingress queues.
show mls qos interface [*interface-id*] [**buffers** | **policers** | **queueing** | **statistics**]	Displays QoS information at the port level, including the buffer allocation, which ports have configured policers, the queueing strategy, and the ingress and egress statistics.
show mls qos maps [**cos-dscp** | **cos-input-q** | **cos-output-q** | **dscp-cos** | **dscp-input-q** | **dscp-mutation** *dscp-mutation-name* | **dscp-output-q** | **ip-prec-dscp** | **policed-dscp**]	Displays QoS mapping information. The DSCP input queue threshold map appears as a matrix. The d1 column specifies the most-significant digit of the DSCP number; the d2 row specifies the least-significant digit in the DSCP number. The intersection of the d1 and the d2 values provides the queue ID and threshold ID; for example, queue 2 and threshold 1 (02-01). The CoS input queue threshold map shows the CoS value in the top row and the corresponding queue ID and threshold ID in the second row; for example, queue 2 and threshold 2 (2-2).
show mls qos maps dscp-to-cos	Verifies your entries.
show mls qos queue-set [*qset-id*]	Displays QoS settings for the egress queues.
show mls qos vlan *vlan-id*	Displays the policy maps attached to the specified SVI.
show policy-map [*policy-map-name* [**class** *class-map-name*]]	Displays QoS policy maps, which define classification criteria for incoming traffic. **Note** Do not use the **show policy-map interface** privileged EXEC command to display classification information for incoming traffic. The **control-plane** and **interface** keywords are not supported, and the statistics shown in the display should be ignored.
show running-config | include rewrite	Displays the DSCP transparency setting.

Configuration Examples for Standard QoS

Configuring the SRR Scheduler: Example

This example shows how to configure the weight ratio of the SRR scheduler running on an egress port. Four queues are used, and the bandwidth ratio allocated for each queue in shared mode is 1/(1+2+3+4), 2/(1+2+3+4), 3/(1+2+3+4), and 4/(1+2+3+4), which is 10 percent, 20 percent, 30 percent, and 40 percent for queues 1, 2, 3, and 4. This means that queue 4 has four times the bandwidth of queue 1, twice the bandwidth of queue 2, and one-and-a-third times the bandwidth of queue 3.

```
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# srr-queue bandwidth share 1 2 3 4
```

Configuring DSCP-Trusted State on a Port: Example

This example shows how to configure a port to the DSCP-trusted state and to modify the DSCP-to-DSCP-mutation map (named *gi0/2-mutation*) so that incoming DSCP values 10 to 13 are mapped to DSCP 30:

```
Switch(config)# mls qos map dscp-mutation gi1/2-mutation 10 11 12 13 to 30
Switch(config)# interface gigabitethernet1/2
Switch(config-if)# mls qos trust dscp
Switch(config-if)# mls qos dscp-mutation gi1/2-mutation
Switch(config-if)# end
```

Allowing ACL Permission for IP Traffic: Examples

This example shows how to allow access for only those hosts on the three specified networks. The wildcard bits apply to the host portions of the network addresses. Any host with a source address that does not match the access list statements is rejected.

```
Switch(config)# access-list 1 permit 192.5.255.0 0.0.0.255
Switch(config)# access-list 1 permit 128.88.0.0 0.0.255.255
Switch(config)# access-list 1 permit 36.0.0.0 0.0.0.255
! (Note: all other access implicitly denied)
```

This example shows how to create an ACL that permits IP traffic from any source to any destination that has the DSCP value set to 32:

```
Switch(config)# access-list 100 permit ip any any dscp 32
```

This example shows how to create an ACL that permits IP traffic from a source host at 10.1.1.1 to a destination host at 10.1.1.2 with a precedence value of 5:

```
Switch(config)# access-list 100 permit ip host 10.1.1.1 host 10.1.1.2 precedence 5
```

This example shows how to create an ACL that permits PIM traffic from any source to a destination group address of 224.0.0.2 with a DSCP set to 32:

```
Switch(config)# access-list 102 permit pim any 224.0.0.2 dscp 32
```

Configuring a Class Map: Examples

This example shows how to configure the class map called *class1*. The *class1* has one match criterion, which is access list 103. It permits traffic from any host to any destination that matches a DSCP value of 10.

```
Switch(config)# access-list 103 permit ip any any dscp 10
Switch(config)# class-map class1
Switch(config-cmap)# match access-group 103
Switch(config-cmap)# end
Switch#
```

This example shows how to create a class map called *class2*, which matches incoming traffic with DSCP values of 10, 11, and 12.

```
Switch(config)# class-map class2
Switch(config-cmap)# match ip dscp 10 11 12
Switch(config-cmap)# end
Switch#
```

This example shows how to create a class map called *class3*, which matches incoming traffic with IP-precedence values of 5, 6, and 7:

```
Switch(config)# class-map class3
Switch(config-cmap)# match ip precedence 5 6 7
Switch(config-cmap)# end
Switch#
```

Creating a Policy Map: Example

This example shows how to create a policy map and attach it to an ingress port. In the configuration, the IP standard ACL permits traffic from network 10.1.0.0. For traffic matching this classification, the DSCP value in the incoming packet is trusted. If the matched traffic exceeds an average traffic rate of 48000 b/s and a normal burst size of 8000 bytes, its DSCP is marked down (based on the policed-DSCP map) and sent:

```
Switch(config)# access-list 1 permit 10.1.0.0 0.0.255.255
Switch(config)# class-map ipclass1
Switch(config-cmap)# match access-group 1
Switch(config-cmap)# exit
Switch(config)# policy-map flow1t
Switch(config-pmap)# class ipclass1
Switch(config-pmap-c)# trust dscp
Switch(config-pmap-c)# police 1000000 8000 exceed-action policed-dscp-transmit
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# service-policy input flow1t
```

Creating a Layer 2 MAC ACL: Example

This example shows how to create a Layer 2 MAC ACL with two permit statements and attach it to an ingress port. The first permit statement allows traffic from the host with MAC address 0001.0000.0001 destined for the host with MAC address 0002.0000.0001. The second permit statement allows only Ethertype XNS-IDP traffic from the host with MAC address 0001.0000.0002 destined for the host with MAC address 0002.0000.0002.

```
Switch(config)# mac access-list extended maclist1
Switch(config-ext-mac)# permit 0001.0000.0001 0.0.0 0002.0000.0001 0.0.0
Switch(config-ext-mac)# permit 0001.0000.0002 0.0.0 0002.0000.0002 0.0.0 xns-idp
Switch(config-ext-mac)# exit
Switch(config)# mac access-list extended maclist2
Switch(config-ext-mac)# permit 0001.0000.0003 0.0.0 0002.0000.0003 0.0.0
Switch(config-ext-mac)# permit 0001.0000.0004 0.0.0 0002.0000.0004 0.0.0 aarp
Switch(config-ext-mac)# exit
Switch(config)# class-map macclass1
Switch(config-cmap)# match access-group maclist1
Switch(config-cmap)# exit
Switch(config)# policy-map macpolicy1
Switch(config-pmap)# class macclass1
Switch(config-pmap-c)# set dscp 63
Switch(config-pmap-c)# exit
Switch(config-pmap)# class macclass2 maclist2
Switch(config-pmap-c)# set dscp 45
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# mls qos trust cos
Switch(config-if)# service-policy input macpolicy1
```

Creating an Aggregate Policer: Example

This example shows how to create an aggregate policer and attach it to multiple classes within a policy map. In the configuration, the IP ACLs permit traffic from network 10.1.0.0 and from host 11.3.1.1. For traffic coming from network 10.1.0.0, the DSCP in the incoming packets is trusted. For traffic coming from host 11.3.1.1, the DSCP in the packet is changed to 56. The traffic rate from the 10.1.0.0 network and from host 11.3.1.1 is policed. If the traffic exceeds an average rate of 48000 b/s and a normal burst size of 8000 bytes, its DSCP is marked down (based on the policed-DSCP map) and sent. The policy map is attached to an ingress port.

```
Switch(config)# access-list 1 permit 10.1.0.0 0.0.255.255
Switch(config)# access-list 2 permit 11.3.1.1
Switch(config)# mls qos aggregate-police transmit1 48000 8000 exceed-action
policed-dscp-transmit
Switch(config)# class-map ipclass1
Switch(config-cmap)# match access-group 1
Switch(config-cmap)# exit
Switch(config)# class-map ipclass2
Switch(config-cmap)# match access-group 2
Switch(config-cmap)# exit
Switch(config)# policy-map aggflow1
Switch(config-pmap)# class ipclass1
Switch(config-pmap-c)# trust dscp
Switch(config-pmap-c)# police aggregate transmit1
Switch(config-pmap-c)# exit
Switch(config-pmap)# class ipclass2
Switch(config-pmap-c)# set dscp 56
Switch(config-pmap-c)# police aggregate transmit1
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# service-policy input aggflow1
Switch(config-if)# exit
```

Configuring COS-to-DSCP Map: Example

This example shows how to modify and display the CoS-to-DSCP map:

```
Switch(config)# mls qos map cos-dscp 10 15 20 25 30 35 40 45
Switch(config)# end
Switch# show mls qos maps cos-dscp

Cos-dscp map:
        cos:   0  1  2  3  4  5  6  7
     --------------------------------
       dscp:   10 15 20 25 30 35 40 45
```

Configuring DSCP Maps: Examples

This example shows how to modify and display the IP-precedence-to-DSCP map:

```
Switch(config)# mls qos map ip-prec-dscp 10 15 20 25 30 35 40 45
Switch(config)# end
Switch# show mls qos maps ip-prec-dscp

IpPrecedence-dscp map:
     ipprec:   0  1  2  3  4  5  6  7
     --------------------------------
       dscp:   10 15 20 25 30 35 40 45
```

This example shows how to map DSCP 50 to 57 to a marked-down DSCP value of 0:

```
Switch(config)# mls qos map policed-dscp 50 51 52 53 54 55 56 57 to 0
Switch(config)# end
Switch# show mls qos maps policed-dscp
Policed-dscp map:
     d1 :  d2 0  1  2  3  4  5  6  7  8  9
     ---------------------------------------
      0 :    00 01 02 03 04 05 06 07 08 09
      1 :    10 11 12 13 14 15 16 17 18 19
      2 :    20 21 22 23 24 25 26 27 28 29
      3 :    30 31 32 33 34 35 36 37 38 39
      4 :    40 41 42 43 44 45 46 47 48 49
      5 :    00 00 00 00 00 00 00 00 58 59
      6 :    60 61 62 63
```

Note In this policed-DSCP map, the marked-down DSCP values are shown in the body of the matrix. The d1 column specifies the most-significant digit of the original DSCP; the d2 row specifies the least-significant digit of the original DSCP. The intersection of the d1 and d2 values provides the marked-down value. For example, an original DSCP value of 53 corresponds to a marked-down DSCP value of 0.

This example shows how to map DSCP values 0, 8, 16, 24, 32, 40, 48, and 50 to CoS value 0 and to display the map:

```
Switch(config)# mls qos map dscp-cos 0 8 16 24 32 40 48 50 to 0
Switch(config)# end
Switch# show mls qos maps dscp-cos
Dscp-cos map:
     d1 :  d2 0  1  2  3  4  5  6  7  8  9
     ---------------------------------------
      0 :    00 00 00 00 00 00 00 00 00 01
      1 :    01 01 01 01 01 01 00 02 02 02
      2 :    02 02 02 02 00 03 03 03 03 03
      3 :    03 03 00 04 04 04 04 04 04 04
      4 :    00 05 05 05 05 05 05 05 00 06
      5 :    00 06 06 06 06 06 07 07 07 07
      6 :    07 07 07 07
```

Note In the above DSCP-to-CoS map, the CoS values are shown in the body of the matrix. The d1 column specifies the most-significant digit of the DSCP; the d2 row specifies the least-significant digit of the DSCP. The intersection of the d1 and d2 values provides the CoS value. For example, in the DSCP-to-CoS map, a DSCP value of 08 corresponds to a CoS value of 0.

This example shows how to define the DSCP-to-DSCP mutation map. All the entries that are not explicitly configured are not modified (remains as specified in the null map):

```
Switch(config)# mls qos map dscp-mutation mutation1 1 2 3 4 5 6 7 to 0
Switch(config)# mls qos map dscp-mutation mutation1 8 9 10 11 12 13 to 10
Switch(config)# mls qos map dscp-mutation mutation1 20 21 22 to 20
Switch(config)# mls qos map dscp-mutation mutation1 30 31 32 33 34 to 30
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# mls qos trust dscp
Switch(config-if)# mls qos dscp-mutation mutation1
Switch(config-if)# end
Switch# show mls qos maps dscp-mutation mutation1
Dscp-dscp mutation map:
   mutation1:
     d1 :  d2 0  1  2  3  4  5  6  7  8  9
     ----------------------------------------
      0 :    00 00 00 00 00 00 00 00 10 10
      1 :    10 10 10 10 14 15 16 17 18 19
      2 :    20 20 20 23 24 25 26 27 28 29
      3 :    30 30 30 30 30 35 36 37 38 39
      4 :    40 41 42 43 44 45 46 47 48 49
      5 :    50 51 52 53 54 55 56 57 58 59
      6 :    60 61 62 63
```

Note In the above DSCP-to-DSCP-mutation map, the mutated values are shown in the body of the matrix. The d1 column specifies the most-significant digit of the original DSCP; the d2 row specifies the least-significant digit of the original DSCP. The intersection of the d1 and d2 values provides the mutated value. For example, a DSCP value of 12 corresponds to a mutated value of 10.

This example shows how to map DSCP values 0 to 6 to ingress queue 1 and to threshold 1 with a drop threshold of 50 percent. It maps DSCP values 20 to 26 to ingress queue 1 and to threshold 2 with a drop threshold of 70 percent:

```
Switch(config)# mls qos srr-queue input dscp-map queue 1 threshold 1 0 1 2 3 4 5 6
Switch(config)# mls qos srr-queue input dscp-map queue 1 threshold 2 20 21 22 23 24 25 26
Switch(config)# mls qos srr-queue input threshold 1 50 70
```

In this example, the DSCP values (0 to 6) are assigned the WTD threshold of 50 percent and will be dropped sooner than the DSCP values (20 to 26) assigned to the WTD threshold of 70 percent.

Configuring an Ingress Queue: Example

This example shows how to allocate 60 percent of the buffer space to ingress queue 1 and 40 percent of the buffer space to ingress queue 2:

```
Switch(config)# mls qos srr-queue input buffers 60 40
```

This example shows how to assign the ingress bandwidth to the queues. Priority queueing is disabled, and the shared bandwidth ratio allocated to queue 1 is 25/(25+75) and to queue 2 is 75/(25+75):

```
Switch(config)# mls qos srr-queue input priority-queue 2 bandwidth 0
Switch(config)# mls qos srr-queue input bandwidth 25 75
```

To return to the default setting, use the **no mls qos srr-queue input priority-queue** *queue-id* global configuration command. To disable priority queueing, set the bandwidth weight to 0, for example, **mls qos srr-queue input priority-queue** *queue-id* **bandwidth 0**.

This example shows how to assign the ingress bandwidths to the queues. Queue 1 is the priority queue with 10 percent of the bandwidth allocated to it. The bandwidth ratios allocated to queues 1 and 2 is 4/(4+4). SRR services queue 1 (the priority queue) first for its configured 10 percent bandwidth. Then SRR equally shares the remaining 90 percent of the bandwidth between queues 1 and 2 by allocating 45 percent to each queue:

```
Switch(config)# mls qos srr-queue input priority-queue 1 bandwidth 10
Switch(config)# mls qos srr-queue input bandwidth 4 4
```

Configuring the Egress Queue: Examples

This example shows how to map a port to queue-set 2. It allocates 40 percent of the buffer space to egress queue 1 and 20 percent to egress queues 2, 3, and 4. It configures the drop thresholds for queue 2 to 40 and 60 percent of the allocated memory, guarantees (reserves) 100 percent of the allocated memory, and configures 200 percent as the maximum memory that this queue can have before packets are dropped:

```
Switch(config)# mls qos queue-set output 2 buffers 40 20 20 20
Switch(config)# mls qos queue-set output 2 threshold 2 40 60 100 200
Switch(config)# interface gigabitethernet1/1
lSwitch(config-if)# queue-set 2
```

This example shows how to map DSCP values 10 and 11 to egress queue 1 and to threshold 2:

```
Switch(config)# mls qos srr-queue output dscp-map queue 1 threshold 2 10 11
```

This example shows how to configure bandwidth shaping on queue 1. Because the weight ratios for queues 2, 3, and 4 are set to 0, these queues operate in shared mode. The bandwidth weight for queue 1 is 1/8, which is 12.5 percent:

```
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# srr-queue bandwidth shape 8 0 0 0
```

This example shows how to enable the egress expedite queue when the SRR weights are configured. The egress expedite queue overrides the configured SRR weights.

```
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# srr-queue bandwidth shape 25 0 0 0
Switch(config-if)# srr-queue bandwidth share 30 20 25 25
Switch(config-if)# priority-queue out
Switch(config-if)# end
```

This example shows how to limit the bandwidth on a port to 80 percent:

```
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# srr-queue bandwidth limit 80
```

When you configure this command to 80 percent, the port is idle 20 percent of the time. The line rate drops to 80 percent of the connected speed, which is 800 Mb/s. These values are not exact because the hardware adjusts the line rate in increments of six.

Creating a Layer 2 MAC ACL: Example

This example shows how to create a Layer 2 MAC ACL with two permit statements. The first statement allows traffic from the host with MAC address 0001.0000.0001 to the host with MAC address 0002.0000.0001. The second statement allows only Ethertype XNS-IDP traffic from the host with MAC address 0001.0000.0002 to the host with MAC address 0002.0000.0002.

```
Switch(config)# mac access-list extended maclist1
```

```
Switch(config-ext-macl)# permit 0001.0000.0001 0.0.0 0002.0000.0001 0.0.0
Switch(config-ext-macl)# permit 0001.0000.0002 0.0.0 0002.0000.0002 0.0.0 xns-idp
! (Note: all other access implicitly denied)
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
Auto-QoS configuration	*Chapter 39, "Configuring Auto-QoS"*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 39

Configuring Auto-QoS

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for Auto-QoS

- When enabling auto-QoS with a Cisco IP phone on a routed port, you must assign a static IP address to the IP phone.
- By default, the CDP is enabled on all ports. For auto-QoS to function properly, do not disable the CDP.

Restrictions for Auto-QoS

- To use this feature, the switch must be running the LAN Base image.
- Connected devices must use Cisco Call Manager Version 4 or later.
- This release supports only Cisco IP SoftPhone Version 1.3(3) or later.
- To take advantage of the auto-QoS defaults, you should enable auto-QoS before you configure other QoS commands. If necessary, you can fine-tune the QoS configuration, but we recommend that you do so only after the auto-QoS configuration is completed. For more information, see the Effects of Auto-QoS on the Configuration, page 39-7.
- Control traffic (such as spanning-tree bridge protocol data units [BPDUs] and routing update packets) received by the switch are subject to all ingress QoS processing.
- You are likely to lose data when you change queue settings; therefore, try to make changes when traffic is at a minimum.

- Auto-QoS configures the switch for VoIP with Cisco IP phones on nonrouted and routed ports. Auto-QoS also configures the switch for VoIP with devices running the Cisco SoftPhone application.
- When a device running Cisco SoftPhone is connected to a nonrouted or routed port, the switch supports only one Cisco SoftPhone application per port.
- Auto-Qos VoIP uses the **priority-queue** interface configuration command for an egress interface. You can also configure a policy-map and trust device on the same interface for Cisco IP phones.
- After auto-QoS is enabled, do not modify a policy map or aggregate policer that includes *AutoQoS* in its name. If you need to modify the policy map or aggregate policer, make a copy of it, and change the copied policy map or policer. To use this new policy map instead of the generated one, remove the generated policy map from the interface, and apply the new policy map to the interface.
- You can enable auto-QoS on static, dynamic-access, voice VLAN access, and trunk ports.

Information About Auto-QoS

This chapter describes how to configure quality of service (QoS) by using automatic QoS (auto-QoS) command on the switch. With QoS, you can provide preferential treatment to certain types of traffic at the expense of others. Without QoS, the switch offers best-effort service to each packet, regardless of the packet contents or size. It sends the packets without any assurance of reliability, delay bounds, or throughput.

You can configure QoS on physical ports and on switch virtual interfaces (SVIs). Other than to apply policy maps, you configure the QoS settings, such as classification, queueing, and scheduling, the same way on physical ports and SVIs. When configuring QoS on a physical port, you apply a nonhierarchical policy map to a port. When configuring QoS on an SVI, you apply a nonhierarchical or a hierarchical policy map.

The switch supports some of the modular QoS CLI (MQC) commands. For more information about the MQC commands, see the “Modular Quality of Service Command-Line Interface Overview” chapter of the *Cisco IOS Quality of Service Solutions Guide*.

Auto-QoS

You can use the auto-QoS feature to simplify the deployment of QoS features. Auto-QoS determines the network design and enables QoS configurations so that the switch can prioritize different traffic flows. It uses the ingress and egress queues instead of using the default (disabled) QoS behavior. The switch offers best-effort service to each packet, regardless of the packet contents or size, and sends it from a single queue.

When you enable auto-QoS, it automatically classifies traffic based on the traffic type and ingress packet label. The switch uses the classification results to choose the appropriate egress queue.

Auto-QoS supports IPv4 and IPv6 traffic when you configure the dual IPv4 and IPv6 SDM template with the **sdm prefer dual ipv4-and-ipv6** global configuration command.

You use auto-QoS commands to identify ports connected to Cisco IP phones and to devices running the Cisco SoftPhone application. You also use the commands to identify ports that receive trusted traffic through an uplink. Auto-QoS then performs these functions:

- Detects the presence or absence of Cisco IP phones
- Configures QoS classification

- Configures egress queues

Generated Auto-QoS Configuration

By default, auto-QoS is disabled on all ports.

When auto-QoS is enabled, it uses the ingress packet label to categorize traffic, to assign packet labels, and to configure the ingress and egress queues as shown in Table 39-1.

Table 39-1 Traffic Types, Packet Labels, and Queues

	VoIP[1] Data Traffic	VoIP Control Traffic	Routing Protocol Traffic	STP BPDU Traffic	Real-Time Video Traffic	All Other Traffic	
DSCP	46	24, 26	48	56	34	–	
CoS	5	3	6	7	4	–	
CoS-to-Ingress Queue Map	2, 3, 4, 5, 6, 7 (queue 2)					0, 1 (queue 1)	
CoS-to-Egress Queue Map	5 (queue 1)	3, 6, 7 (queue 2)			4 (queue 3)	2 (queue 3)	0, 1 (queue 4)

1. VoIP = voice over IP

Table 39-2 shows the generated auto-QoS configuration for the ingress queues.

Table 39-2 Auto-QoS Configuration for the Ingress Queues

Ingress Queue	Queue Number	CoS-to-Queue Map	Queue Weight (Bandwidth)	Queue (Buffer) Size
SRR shared	1	0, 1	81 percent	67 percent
Priority	2	2, 3, 4, 5, 6, 7	19 percent	33 percent

Table 39-3 shows the generated auto-QoS configuration for the egress queues.

Table 39-3 Auto-QoS Configuration for the Egress Queues

Egress Queue	Queue Number	CoS-to-Queue Map	Queue Weight (Bandwidth)	Queue (Buffer) Size for Gigabit-Capable Ports	Queue (Buffer) Size for 10/100 Ethernet Ports
Priority	1	5	up to100 percent	16 percent	10 percent
SRR shared	2	3, 6, 7	10 percent	6 percent	10 percent
SRR shared	3	2, 4	60 percent	17 percent	26 percent
SRR shared	4	0, 1	20 percent	61 percent	54 percent

When you enable the auto-QoS feature on the first port, these automatic actions occur:

- QoS is globally enabled (**mls qos** global configuration command), and other global configuration commands are added.
- When you enter the **auto qos voip cisco-phone** interface configuration command on a port at the edge of the network that is connected to a Cisco IP phone, the switch enables the trusted boundary feature. The switch uses the Cisco Discovery Protocol (CDP) to detect the presence or absence of a Cisco IP phone. When a Cisco IP phone is detected, the ingress classification on the port is set to trust the QoS label received in the packet. The switch also uses policing to determine whether a packet is in or out of profile and to specify the action on the packet. If the packet does not have a DSCP value of 24, 26, or 46 or is out of profile, the switch changes the DSCP value to 0. When a Cisco IP phone is absent, the ingress classification is set to not trust the QoS label in the packet. The switch configures ingress and egress queues on the port according to the settings in Table 39-2 and Table 39-3. The policing is applied to those traffic matching the policy-map classification before the switch enables the trust boundary feature.
- When you enter the **auto qos voip cisco-softphone** interface configuration command on a port at the edge of the network that is connected to a device running the Cisco SoftPhone, the switch uses policing to determine whether a packet is in or out of profile and to specify the action on the packet. If the packet does not have a DSCP value of 24, 26, or 46 or is out of profile, the switch changes the DSCP value to 0. The switch configures ingress and egress queues on the port according to the settings in Table 39-2 and Table 39-3.
- When you enter the **auto qos voip trust** interface configuration command on a port connected to the interior of the network, the switch trusts the CoS value for nonrouted ports or the DSCP value for routed ports in ingress packets (the assumption is that traffic has already been classified by other edge devices). The switch configures the ingress and egress queues on the port according to the settings in Table 39-2 and Table 39-3.

  For information about the trusted boundary feature, see the "Configuring a Trusted Boundary to Ensure Port Security" section on page 38-34.

When you enable auto-QoS by using the **auto qos voip cisco-phone**, the **auto qos voip cisco-softphone**, or the **auto qos voip trust** interface configuration command, the switch automatically generates a QoS configuration based on the traffic type and ingress packet label and applies the commands listed in Table 39-4 to the port.

Table 39-4 Generated Auto-QoS Configuration

Description	Automatically Generated Command
The switch automatically enables standard QoS and configures the CoS-to-DSCP map (maps CoS values in incoming packets to a DSCP value).	`Switch(config)# mls qos` `Switch(config)# mls qos map cos-dscp 0 8 16 26 32 46 48 56`

The switch automatically maps CoS values to an ingress queue and to a threshold ID.	`Switch(config)# no mls qos srr-queue input cos-map`
	`Switch(config)# mls qos srr-queue input cos-map queue 1 threshold 3 0`
	`Switch(config)# mls qos srr-queue input cos-map queue 1 threshold 2 1`
	`Switch(config)# mls qos srr-queue input cos-map queue 2 threshold 1 2`
	`Switch(config)# mls qos srr-queue input cos-map queue 2 threshold 2 4 6 7`
	`Switch(config)# mls qos srr-queue input cos-map queue 2 threshold 3 3 5`

Table 39-4 Generated Auto-QoS Configuration (continued)

Description	Automatically Generated Command
The switch automatically maps CoS values to an egress queue and to a threshold ID.	Switch(config)# **no mls qos srr-queue output cos-map** Switch(config)# **mls qos srr-queue output cos-map queue 1 threshold 3 5** Switch(config)# **mls qos srr-queue output cos-map queue 2 threshold 3 3 6 7** Switch(config)# **mls qos srr-queue output cos-map queue 3 threshold 3 2 4** Switch(config)# **mls qos srr-queue output cos-map queue 4 threshold 2 1** Switch(config)# **mls qos srr-queue output cos-map queue 4 threshold 3 0**
The switch automatically maps DSCP values to an ingress queue and to a threshold ID.	Switch(config)# **no mls qos srr-queue input dscp-map** Switch(config)# **mls qos srr-queue input dscp-map queue 1 threshold 2 9 10 11 12 13 14 15** Switch(config)# **mls qos srr-queue input dscp-map queue 1 threshold 3 0 1 2 3 4 5 6 7** Switch(config)# **mls qos srr-queue input dscp-map queue 1 threshold 3 32** Switch(config)# **mls qos srr-queue input dscp-map queue 2 threshold 1 16 17 18 19 20 21 22 23** Switch(config)# **mls qos srr-queue input dscp-map queue 2 threshold 2 33 34 35 36 37 38 39 48** Switch(config)# **mls qos srr-queue input dscp-map queue 2 threshold 2 49 50 51 52 53 54 55 56** Switch(config)# **mls qos srr-queue input dscp-map queue 2 threshold 2 57 58 59 60 61 62 63** Switch(config)# **mls qos srr-queue input dscp-map queue 2 threshold 3 24 25 26 27 28 29 30 31** Switch(config)# **mls qos srr-queue input dscp-map queue 2 threshold 3 40 41 42 43 44 45 46 47**
The switch automatically maps DSCP values to an egress queue and to a threshold ID.	Switch(config)# **no mls qos srr-queue output dscp-map** Switch(config)# **mls qos srr-queue output dscp-map queue 1 threshold 3 40 41 42 43 44 45 46 47** Switch(config)# **mls qos srr-queue output dscp-map queue 2 threshold 3 24 25 26 27 28 29 30 31** Switch(config)# **mls qos srr-queue output dscp-map queue 2 threshold 3 48 49 50 51 52 53 54 55** Switch(config)# **mls qos srr-queue output dscp-map queue 2 threshold 3 56 57 58 59 60 61 62 63** Switch(config)# **mls qos srr-queue output dscp-map queue 3 threshold 3 16 17 18 19 20 21 22 23** Switch(config)# **mls qos srr-queue output dscp-map queue 3 threshold 3 32 33 34 35 36 37 38 39** Switch(config)# **mls qos srr-queue output dscp-map queue 4 threshold 1 8** Switch(config)# **mls qos srr-queue output dscp-map queue 4 threshold 2 9 10 11 12 13 14 15** Switch(config)# **mls qos srr-queue output dscp-map queue 4 threshold 3 0 1 2 3 4 5 6 7**

Table 39-4 Generated Auto-QoS Configuration (continued)

Description	Automatically Generated Command

The switch automatically sets up the ingress queues, with queue 2 as the priority queue and queue 1 in shared mode. The switch also configures the bandwidth and buffer size for the ingress queues.	`Switch(config)# no mls qos srr-queue input priority-queue 1`
	`Switch(config)# no mls qos srr-queue input priority-queue 2`
	`Switch(config)# mls qos srr-queue input bandwidth 90 10`
	`Switch(config)# mls qos srr-queue input threshold 1 8 16`
	`Switch(config)# mls qos srr-queue input threshold 2 34 66`
	`Switch(config)# mls qos srr-queue input buffers 67 33`

The switch automatically configures the egress queue buffer sizes. It configures the bandwidth and the SRR mode (shaped or shared) on the egress queues mapped to the port.	`Switch(config)# mls qos queue-set output 1 threshold 1 138 138 92 138`
	`Switch(config)# mls qos queue-set output 1 threshold 2 138 138 92 400`
	`Switch(config)# mls qos queue-set output 1 threshold 3 36 77 100 318`
	`Switch(config)# mls qos queue-set output 1 threshold 4 20 50 67 400`
	`Switch(config)# mls qos queue-set output 2 threshold 1 149 149 100 149`
	`Switch(config)# mls qos queue-set output 2 threshold 2 118 118 100 235`
	`Switch(config)# mls qos queue-set output 2 threshold 3 41 68 100 272`
	`Switch(config)# mls qos queue-set output 2 threshold 4 42 72 100 242`
	`Switch(config)# mls qos queue-set output 1 buffers 10 10 26 54`
	`Switch(config)# mls qos queue-set output 2 buffers 16 6 17 61`
	`Switch(config-if)# priority-que out`
	`Switch(config-if)# srr-queue bandwidth share 10 10 60 20`

If you entered the **auto qos voip trust** command, the switch automatically sets the ingress classification to trust the CoS value received in the packet on a nonrouted port by using the **mls qos trust cos** command or to trust the DSCP value received in the packet on a routed port by using the **mls qos trust dscp** command.	`Switch(config-if)# mls qos trust cos` `Switch(config-if)# mls qos trust dscp`
If you entered the **auto qos voip cisco-phone** command, the switch automatically enables the trusted boundary feature, which uses the CDP to detect the presence or absence of a Cisco IP phone.	`Switch(config-if)# mls qos trust device cisco-phone`

Table 39-4 Generated Auto-QoS Configuration (continued)

Description	Automatically Generated Command

If you entered the **auto qos voip cisco-softphone** command, the switch automatically creates class maps and policy maps.	`Switch(config)# mls qos map policed-dscp 24 26 46 to 0`
	`Switch(config)# class-map match-all AutoQoS-VoIP-RTP-Trust`
	`Switch(config-cmap)# match ip dscp ef`
	`Switch(config)# class-map match-all AutoQoS-VoIP-Control-Trust`
	`Switch(config-cmap)# match ip dscp cs3 af31`
	`Switch(config)# policy-map AutoQoS-Police-SoftPhone`
	`Switch(config-pmap)# class AutoQoS-VoIP-RTP-Trust`
	`Switch(config-pmap-c)# set dscp ef`
	`Switch(config-pmap-c)# police 320000 8000 exceed-action policed-dscp-transmit`
	`Switch(config-pmap)# class AutoQoS-VoIP-Control-Trust`
	`Switch(config-pmap-c)# set dscp cs3`
	`Switch(config-pmap-c)# police 32000 8000 exceed-action policed-dscp-transmit`

After creating the class maps and policy maps, the switch automatically applies the policy map called *AutoQoS-Police-SoftPhone* to an ingress interface on which auto-QoS with the Cisco SoftPhone feature is enabled.	`Switch(config-if)# service-policy input AutoQoS-Police-SoftPhone`

If you entered the **auto qos voip cisco-phone** command, the switch automatically creates class maps and policy maps.	`witch(config)# mls qos map policed-dscp 24 26 46 to 0`
	`Switch(config)# class-map match-all AutoQoS-VoIP-RTP-Trust`
	`Switch(config-cmap)# match ip dscp ef`
	`Switch(config)# class-map match-all AutoQoS-VoIP-Control-Trust`
	`Switch(config-cmap)# match ip dscp cs3 af31`
	`Switch(config)# policy-map AutoQoS-Police-CiscoPhone`
	`Switch(config-pmap)# class AutoQoS-VoIP-RTP-Trust`
	`Switch(config-pmap-c)# set dscp ef`
	`Switch(config-pmap-c)# police 320000 8000 exceed-action policed-dscp-transmit`
	`Switch(config-pmap)# class AutoQoS-VoIP-Control-Trust`
	`Switch(config-pmap-c)# set dscp cs3`
	`Switch(config-pmap-c)# police 32000 8000 exceed-action policed-dscp-transmit`

After creating the class maps and policy maps, the switch automatically applies the policy map named *AutoQoS-Police-CiscoPhone* to an ingress interface on which auto-QoS with the Cisco IP phone feature is enabled.	`Switch(config-if)# service-policy input AutoQoS-Police-CiscoPhone`

Effects of Auto-QoS on the Configuration

When auto-QoS is enabled, the **auto qos voip** interface configuration command and the generated configuration are added to the running configuration.

The switch applies the auto-QoS-generated commands as if the commands were entered from the CLI. An existing user configuration can cause the application of the generated commands to fail or to be overridden by the generated commands. These actions occur without warning. If all the generated

commands are successfully applied, any user-entered configuration that was not overridden remains in the running configuration. Any user-entered configuration that was overridden can be retrieved by reloading the switch without saving the current configuration to memory. If the generated commands fail to be applied, the previous running configuration is restored.

To display the QoS commands that are automatically generated when auto-QoS is enabled or disabled, enter the **debug auto qos** privileged EXEC command *before* enabling auto-QoS. For more information, see the **debug autoqos** command in the command reference for this release.

To disable auto-QoS on a port, use the **no auto qos voip** interface configuration command. Only the auto-QoS-generated interface configuration commands for this port are removed. If this is the last port on which auto-QoS is enabled and you enter the **no auto qos voip** command, auto-QoS is considered disabled even though the auto-QoS-generated global configuration commands remain (to avoid disrupting traffic on other ports affected by the global configuration).

You can use the **no mls qos** global configuration command to disable the auto-QoS-generated global configuration commands. With QoS disabled, there is no concept of trusted or untrusted ports because the packets are not modified (the CoS, DSCP, and IP precedence values in the packet are not changed). Traffic is switched in pass-through mode (packets are switched without any rewrites and classified as best effort without any policing).

How to Configure Auto-QoS

Enabling Auto-QoS for VoIP

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the port that is connected to a Cisco IP phone, the port that is connected to a device running the Cisco SoftPhone feature, or the uplink port that is connected to another trusted switch or router in the interior of the network, and enter interface configuration mode.
Step 3	**auto qos voip** {**cisco-phone** | **cisco-softphone** | **trust**}	Enables auto-QoS. • **cisco-phone**—Specifies the port is connected to a Cisco IP phone, the QoS labels of incoming packets are trusted only when the telephone is detected. • **cisco-softphone**—Specifies the port is connected to a device running the Cisco SoftPhone feature. • **trust**—Specifies the uplink port is connected to a trusted switch or router, and the VoIP traffic classification in the ingress packet is trusted.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring QoS to Prioritize VoIP Traffic

This task explains how to configure the switch at the edge of the QoS domain to prioritize the VoIP traffic over all other traffic:

	Command	Purpose
Step 1	**debug auto qos**	Enables debugging for auto-QoS. When debugging is enabled, the switch displays the QoS configuration that is automatically generated when auto-QoS is enabled.
Step 2	**configure terminal**	Enters global configuration mode.
Step 3	**cdp enable**	Enable CDP globally. By default, it is enabled.
Step 4	**interface** *interface-id*	Specifies the switch port connected to the Cisco IP phone, and enters interface configuration mode.
Step 5	**auto qos voip cisco-phone**	Enables auto-QoS on the port, and specifies that the port is connected to a Cisco IP phone. The QoS labels of incoming packets are trusted only when the Cisco IP phone is detected.
Step 6	**exit**	Returns to global configuration mode.
Step 7		Repeat Steps 4 to 6 for as many ports as are connected to the Cisco IP phone.
Step 8	**interface** *interface-id*	Specifies the switch port identified as connected to a trusted switch or router, and enters interface configuration mode. See Figure 39-1.
Step 9	**auto qos voip trust**	Enables auto-QoS on the port, and specifies that the port is connected to a trusted router or switch.
Step 10	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining Auto-QoS

Command	Purpose
show auto qos [**interface** [*interface-id*]]	Displays the QoS commands entered on the interfaces on which auto-QoS is enabled.
show mls qos	Displays global QoS configuration information.
show mls qos interface [*interface-id*] [**buffers** | **queueing**]	Displays QoS information at the port level.
show mls qos maps [**cos-dscp** | **cos-input-q** | **cos-output-q** | **dscp-cos** | **dscp-input-q** | **dscp-mutation** | **dscp-output-q** | **ip-prec-dscp** | **policed-dscp**]	Displays QoS mapping information. During classification, QoS uses the mapping tables to represent the priority of the traffic and to derive a corresponding CoS or DSCP value from the received CoS, DSCP, or IP precedence value.
show mls qos input-queue	Displays QoS settings for the ingress queues.
show running-config	Displays the current operating configuration, including defined macros.

Configuration Examples for Auto-QoS

Auto-QoS Network: Example

This is an illustrated example that shows you how to implement auto-QoS in a network in which the VoIP traffic is prioritized over all other traffic. Auto-QoS is enabled on the switches in the wiring closets at the edge of the QoS domain.

For optimum QoS performance, enable auto-QoS on all the devices in the network.

Figure 39-1 Auto-QoS Configuration Example Network

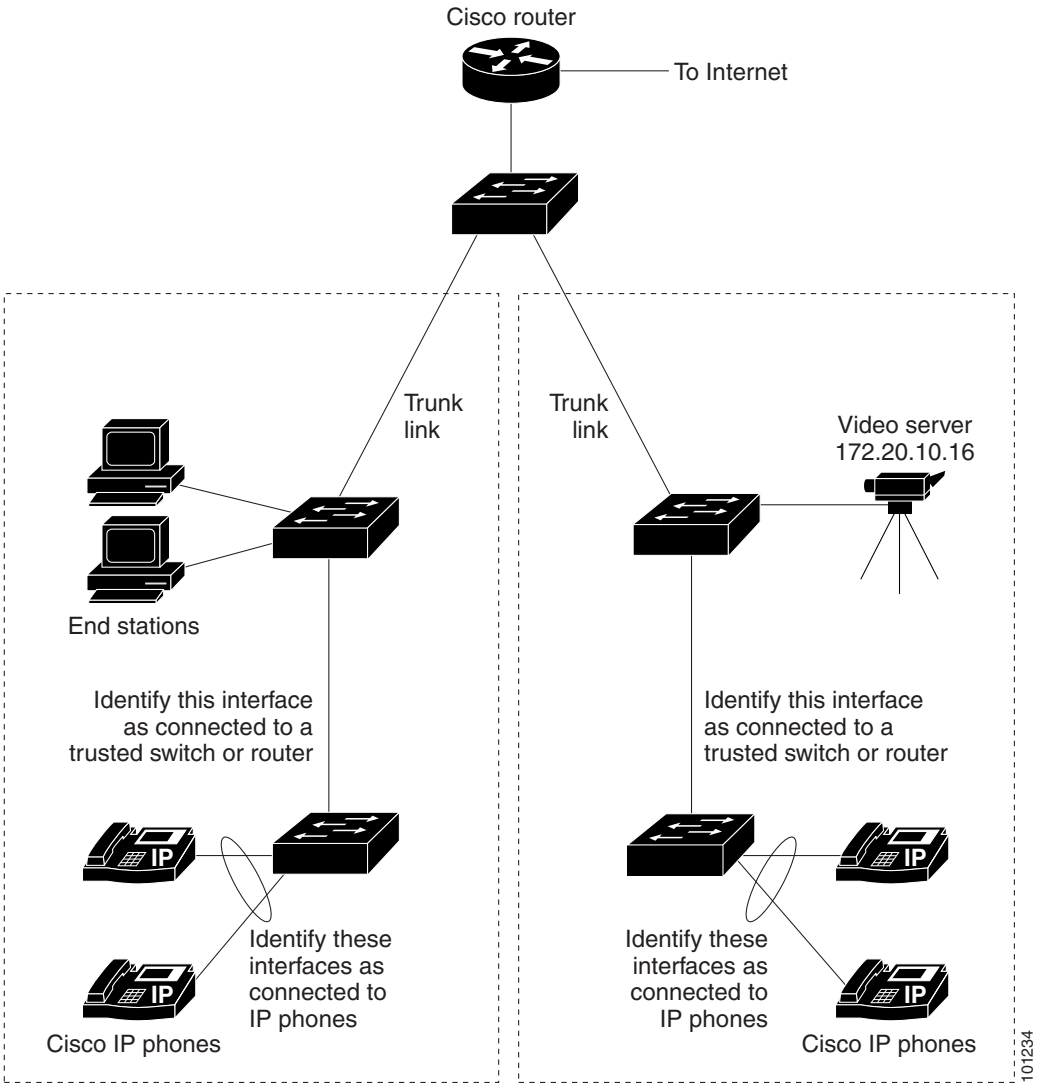

Enabling Auto-QoS VOIP Trust: Example

This example shows how to enable auto-QoS and to trust the QoS labels received in incoming packets when the switch or router connected to a port is a trusted device:

```
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# auto qos voip trust
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
Standard QoS	*Chapter 38, "Configuring Standard QoS"*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 40

Configuring EtherChannels

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Restrictions for Configuring EtherChannels

- To use this feature, the switch must be running the LAN Base image.
- Port channel is supported in only the LAN Base image.

Information About Configuring EtherChannels

This chapter describes how to configure EtherChannels on the switch. EtherChannel provides fault-tolerant high-speed links between switches, routers, and servers. You can use it to increase the bandwidth between the wiring closets and the data center, and you can deploy it anywhere in the network where bottlenecks are likely to occur. EtherChannel provides automatic recovery for the loss of a link by redistributing the load across the remaining links. If a link fails, EtherChannel redirects traffic from the failed link to the remaining links in the channel without intervention. This chapter also describes how to configure link-state tracking.

EtherChannels

An EtherChannel consists of individual Fast Ethernet or Gigabit Ethernet links bundled into a single logical link as shown in Figure 40-1.

Figure 40-1 Typical EtherChannel Configuration

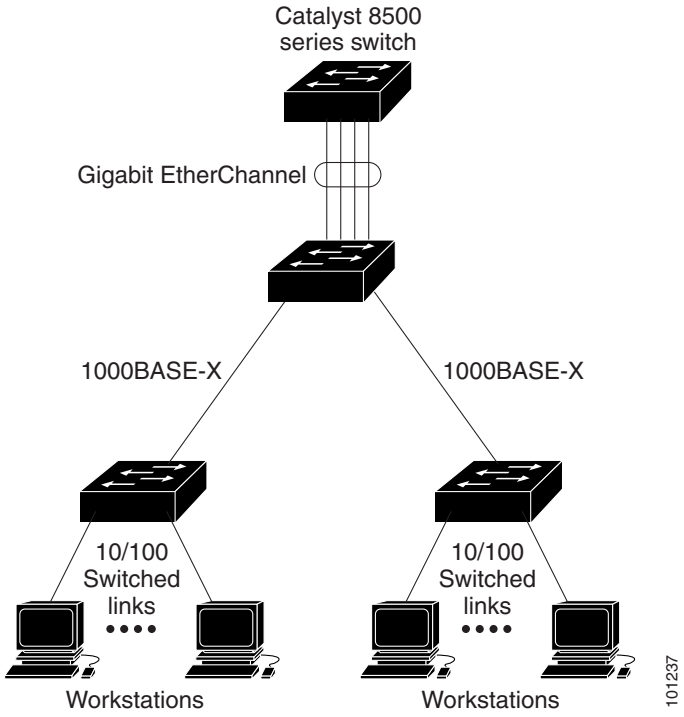

The EtherChannel provides full-duplex bandwidth up to 800 Mb/s (Fast EtherChannel) or 2 Gb/s (Gigabit EtherChannel) between your switch and another switch or host. Each EtherChannel can consist of up to eight compatibly configured Ethernet ports.

The number of EtherChannels is limited to six. For more information, see the "EtherChannel Configuration Guidelines" section on page 40-10.

You can configure an EtherChannel in one of these modes: Port Aggregation Protocol (PAgP), Link Aggregation Control Protocol (LACP), or On. Configure both ends of the EtherChannel in the same mode:

- When you configure one end of an EtherChannel in either PAgP or LACP mode, the system negotiates with the other end of the channel to determine which ports should become active. Incompatible ports are put into an independent state and continue to carry data traffic as would any other single link. The port configuration does not change, but the port does not participate in the EtherChannel.
- When you configure an EtherChannel in the **on** mode, no negotiations take place. The switch forces all compatible ports to become active in the EtherChannel. The other end of the channel (on the other switch) must also be configured in the **on** mode; otherwise, packet loss can occur.

If a link within an EtherChannel fails, traffic previously carried over that failed link moves to the remaining links within the EtherChannel. If traps are enabled on the switch, a trap is sent for a failure that identifies the switch, the EtherChannel, and the failed link. Inbound broadcast and multicast packets on one link in an EtherChannel are blocked from returning on any other link of the EtherChannel.

Port-Channel Interfaces

When you create an EtherChannel, a port-channel logical interface is involved:

- With Layer 2 ports, use the **channel-group** interface configuration command to dynamically create the port-channel logical interface.

  You also can use the **interface port-channel** *port-channel-number* global configuration command to manually create the port-channel logical interface, but then you must use the **channel-group** *channel-group-number* command to bind the logical interface to a physical port. The *channel-group-number* can be the same as the *port-channel-number,* or you can use a new number. If you use a new number, the **channel-group** command dynamically creates a new port channel.

- With Layer 3 ports, you should manually create the logical interface by using the **interface port-channel** global configuration command followed by the **no switchport** interface configuration command. Then you manually assign an interface to the EtherChannel by using the **channel-group** interface configuration command.

For both Layer 2 and Layer 3 ports, the **channel-group** command binds the physical port and the logical interface together as shown in Figure 40-2.

Each EtherChannel has a port-channel logical interface numbered from 1 to 6. This port-channel interface number corresponds to the one specified with the **channel-group** interface configuration command.

Figure 40-2 Relationship of Physical Ports, Logical Port Channels, and Channel Groups

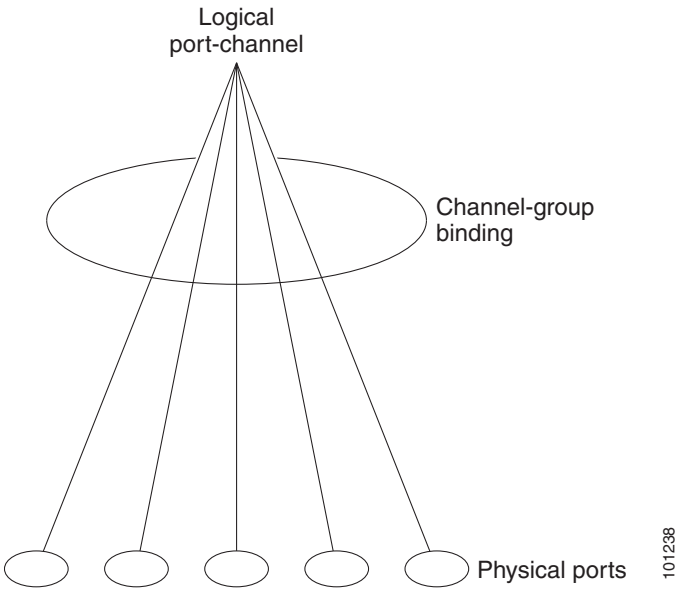

After you configure an EtherChannel, configuration changes applied to the port-channel interface apply to all the physical ports assigned to the port-channel interface. Configuration changes applied to the physical port affect only the port where you apply the configuration. To change the parameters of all ports in an EtherChannel, apply configuration commands to the port-channel interface, for example, spanning-tree commands or commands to configure a Layer 2 EtherChannel as a trunk.

Port Aggregation Protocol

The Port Aggregation Protocol (PAgP) is a Cisco-proprietary protocol that can be run only on Cisco switches and on those switches licensed by vendors to support PAgP. PAgP facilitates the automatic creation of EtherChannels by exchanging PAgP packets between Ethernet ports.

By using PAgP, the switch learns the identity of partners capable of supporting PAgP and the capabilities of each port. It then dynamically groups similarly configured ports into a single logical link (channel or aggregate port). Similarly configured ports are grouped based on hardware, administrative, and port parameter constraints. For example, PAgP groups the ports with the same speed, duplex mode, native VLAN, VLAN range, and trunking status and type. After grouping the links into an EtherChannel, PAgP adds the group to the spanning tree as a single switch port.

PAgP Modes

Table 40-1 shows the user-configurable EtherChannel PAgP modes for the **channel-group** interface configuration command.

Table 40-1 EtherChannel PAgP Modes

Mode	Description
auto	Places a port into a passive negotiating state, in which the port responds to PAgP packets it receives but does not start PAgP packet negotiation. This setting minimizes the transmission of PAgP packets.
desirable	Places a port into an active negotiating state, in which the port starts negotiations with other ports by sending PAgP packets.

Switch ports exchange PAgP packets only with partner ports configured in the **auto** or **desirable** modes. Ports configured in the **on** mode do not exchange PAgP packets.

Both the **auto** and **desirable** modes enable ports to negotiate with partner ports to form an EtherChannel based on criteria such as port speed and, for Layer 2 EtherChannels, trunking state and VLAN numbers.

Ports can form an EtherChannel when they are in different PAgP modes as long as the modes are compatible. For example:

- A port in the **desirable** mode can form an EtherChannel with another port that is in the **desirable** or **auto** mode.
- A port in the **auto** mode can form an EtherChannel with another port in the **desirable** mode.

A port in the **auto** mode cannot form an EtherChannel with another port that is also in the **auto** mode because neither port starts PAgP negotiation.

If your switch is connected to a partner that is PAgP-capable, you can configure the switch port for nonsilent operation by using the **non-silent** keyword. If you do not Specifies **non-silent** with the **auto** or **desirable** mode, silent mode is assumed.

Use the silent mode when the switch is connected to a device that is not PAgP-capable and seldom, if ever, sends packets. An example of a silent partner is a file server or a packet analyzer that is not generating traffic. In this case, running PAgP on a physical port connected to a silent partner prevents that switch port from ever becoming operational. However, the silent setting allows PAgP to operate, to attach the port to a channel group, and to use the port for transmission.

PAgP Learn Method and Priority

Network devices are classified as PAgP physical learners or aggregate-port learners. A device is a physical learner if it learns addresses by physical ports and directs transmissions based on that knowledge. A device is an aggregate-port learner if it learns addresses by aggregate (logical) ports. The learn method must be configured the same at both ends of the link.

When a device and its partner are both aggregate-port learners, they learn the address on the logical port-channel. The device sends packets to the source by using any of the ports in the EtherChannel. With aggregate-port learning, it is not important on which physical port the packet arrives.

PAgP cannot automatically detect when the partner device is a physical learner and when the local device is an aggregate-port learner. Therefore, you must manually set the learning method on the local device to learn addresses by physical ports. You also must set the load-distribution method to source-based distribution, so that any given source MAC address is always sent on the same physical port.

You also can configure a single port within the group for all transmissions and use other ports for hot standby. The unused ports in the group can be swapped into operation in just a few seconds if the selected single port loses hardware-signal detection. You can configure which port is always selected for packet transmission by changing its priority with the **pagp port-priority** interface configuration command. The higher the priority, the more likely that the port will be selected.

Note The switch supports address learning only on aggregate ports even though the **physical-port** keyword is provided in the CLI. The **pagp learn-method** command and the **pagp port-priority** command have no effect on the switch hardware, but they are required for PAgP interoperability with devices that only support address learning by physical ports.

When the link partner of the switch is a physical learner (such as a Catalyst 1900 series switch), we recommend that you configure the switch as a physical-port learner by using the **pagp learn-method physical-port** interface configuration command. Set the load-distribution method based on the source MAC address by using the **port-channel load-balance src-mac** global configuration command. The switch then sends packets to the Catalyst 1900 switch using the same port in the EtherChannel from which it learned the source address. Only use the **pagp learn-method** command in this situation.

PAgP Interaction with Virtual Switches and Dual-Active Detection

A virtual switch can be two or more Catalyst 6500 core switches connected by virtual switch links (VSLs) that carry control and data traffic between them. One of the switches is in active mode. The others are in standby mode. For redundancy, remote switches, are connected to the virtual switch by remote satellite links (RSLs).

If the VSL between two switches fails, one switch does not know the status of the other. Both switches could change to the active mode, causing a *dual-active situation* in the network with duplicate configurations (including duplicate IP addresses and bridge identifiers). The network might go down.

To prevent a dual-active situation, the core switches send PAgP protocol data units (PDUs) through the RSLs to the remote switches. The PAgP PDUs identify the active switch, and the remote switches forward the PDUs to core switches so that the core switches are in sync. If the active switch fails or resets, the standby switch takes over as the active switch. If the VSL goes down, one core switch knows the status of the other and does not change state.

PAgP Interaction with Other Features

The Dynamic Trunking Protocol (DTP) and the Cisco Discovery Protocol (CDP) send and receive packets over the physical ports in the EtherChannel. Trunk ports send and receive PAgP protocol data units (PDUs) on the lowest numbered VLAN.

In Layer 2 EtherChannels, the first port in the channel that comes up provides its MAC address to the EtherChannel. If this port is removed from the bundle, one of the remaining ports in the bundle provides its MAC address to the EtherChannel.

PAgP sends and receives PAgP PDUs only from ports that are up and have PAgP enabled for the auto or desirable mode.

Link Aggregation Control Protocol

The LACP is defined in IEEE 802.3ad and enables Cisco switches to manage Ethernet channels between switches that conform to the IEEE 802.3ad protocol. LACP facilitates the automatic creation of EtherChannels by exchanging LACP packets between Ethernet ports.

By using LACP, the switch learns the identity of partners capable of supporting LACP and the capabilities of each port. It then dynamically groups similarly configured ports into a single logical link (channel or aggregate port). Similarly configured ports are grouped based on hardware, administrative, and port parameter constraints. For example, LACP groups the ports with the same speed, duplex mode, native VLAN, VLAN range, and trunking status and type. After grouping the links into an EtherChannel, LACP adds the group to the spanning tree as a single switch port.

LACP Modes

Table 40-2 shows the user-configurable EtherChannel LACP modes for the **channel-group** interface configuration command.

Table 40-2 EtherChannel LACP Modes

Mode	Description
active	Places a port into an active negotiating state in which the port starts negotiations with other ports by sending LACP packets.
passive	Places a port into a passive negotiating state in which the port responds to LACP packets that it receives, but does not start LACP packet negotiation. This setting minimizes the transmission of LACP packets.

Both the **active** and **passive LACP** modes enable ports to negotiate with partner ports to an EtherChannel based on criteria such as port speed and, for Layer 2 EtherChannels, trunking state and VLAN numbers.

Ports can form an EtherChannel when they are in different LACP modes as long as the modes are compatible. For example:

- A port in the **active** mode can form an EtherChannel with another port that is in the **active** or **passive** mode.
- A port in the **passive** mode cannot form an EtherChannel with another port that is also in the **passive** mode because neither port starts LACP negotiation.

LACP Hot-Standby Ports

When enabled, LACP tries to configure the maximum number of LACP-compatible ports in a channel, up to a maximum of 16 ports. Only eight LACP links can be active at one time. The software places any additional links in a hot-standby mode. If one of the active links becomes inactive, a link that is in the hot-standby mode becomes active in its place.

If you configure more than eight links for an EtherChannel group, the software automatically decides which of the hot-standby ports to make active based on the LACP priority. To every link between systems that operate LACP, the software assigns a unique priority made up of these elements (in priority order):

- LACP system priority
- System ID (the switch MAC address)
- LACP port priority
- Port number

In priority comparisons, numerically lower values have higher priority. The priority decides which ports should be put in standby mode when there is a hardware limitation that prevents all compatible ports from aggregating.

Determining which ports are active and which are hot standby is a two-step procedure. First the system with a numerically lower system priority and system-id is placed in charge of the decision. Next, that system decides which ports are active and which are hot standby, based on its values for port priority and port number. The port-priority and port-number values for the other system are not used.

You can change the default values of the LACP system priority and the LACP port priority to affect how the software selects active and standby links.

By default, all ports use the same port priority. If the local system has a lower value for the system priority and the system ID than the remote system, you can affect which of the hot-standby links become active first by changing the port priority of LACP EtherChannel ports to a lower value than the default**.** The hot-standby ports that have lower port numbers become active in the channel first. You can use the **show etherchannel summary** privileged EXEC command to see which ports are in the hot-standby mode (denoted with an *H* port-state flag).

If LACP is not able to aggregate all the ports that are compatible (for example, the remote system might have more restrictive hardware limitations), all the ports that cannot be actively included in the EtherChannel are put in the hot-standby state and are used only if one of the channeled ports fails.

LACP Interaction with Other Features

The DTP and the CDP send and receive packets over the physical ports in the EtherChannel. Trunk ports send and receive LACP PDUs on the lowest numbered VLAN.

In Layer 2 EtherChannels, the first port in the channel that comes up provides its MAC address to the EtherChannel. If this port is removed from the bundle, one of the remaining ports in the bundle provides its MAC address to the EtherChannel.

LACP sends and receives LACP PDUs only from ports that are up and have LACP enabled for the active or passive mode.

EtherChannel On Mode

EtherChannel **on** mode can be used to manually configure an EtherChannel. The **on** mode forces a port to join an EtherChannel without negotiations. The **on** mode can be useful if the remote device does not support PAgP or LACP. In the **on** mode, a usable EtherChannel exists only when the switches at both ends of the link are configured in the **on** mode.

Ports that are configured in the **on** mode in the same channel group must have compatible port characteristics, such as speed and duplex. Ports that are not compatible are suspended, even though they are configured in the **on** mode.

Caution You should use care when using the **on** mode. This is a manual configuration, and ports on both ends of the EtherChannel must have the same configuration. If the group is misconfigured, packet loss or spanning-tree loops can occur.

Load Balancing and Forwarding Methods

EtherChannel balances the traffic load across the links in a channel by reducing part of the binary pattern formed from the addresses in the frame to a numerical value that selects one of the links in the channel. EtherChannel load balancing can use MAC addresses or IP addresses, source or destination addresses, or both source and destination addresses. The selected mode applies to all EtherChannels configured on the switch. You configure the load balancing and forwarding method by using the **port-channel load-balance** global configuration command.

With source-MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on the source-MAC address of the incoming packet. Therefore, to provide load balancing, packets from different hosts use different ports in the channel, but packets from the same host use the same port in the channel.

With destination-MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on the destination host's MAC address of the incoming packet. Therefore, packets to the same destination are forwarded over the same port, and packets to a different destination are sent on a different port in the channel.

With source-and-destination MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on both the source and destination MAC addresses. This forwarding method, a combination source-MAC and destination-MAC address forwarding methods of load distribution, can be used if it is not clear whether source-MAC or destination-MAC address forwarding is better suited on a particular switch. With source-and-destination MAC-address forwarding, packets sent from host A to host B, host A to host C, and host C to host B could all use different ports in the channel.

With source-IP address-based forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the EtherChannel based on the source-IP address of the incoming packet. Therefore, to provide load-balancing, packets from different IP addresses use different ports in the channel, but packets from the same IP address use the same port in the channel.

With destination-IP address-based forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the EtherChannel based on the destination-IP address of the incoming packet. Therefore, to provide load-balancing, packets from the same IP source address sent to different IP destination addresses could be sent on different ports in the channel. But packets sent from different source IP addresses to the same destination IP address are always sent on the same port in the channel.

With source-and-destination IP address-based forwarding, packets are sent to an EtherChannel and distributed across the EtherChannel ports, based on both the source and destination IP addresses of the incoming packet. This forwarding method, a combination of source-IP and destination-IP address-based forwarding, can be used if it is not clear whether source-IP or destination-IP address-based forwarding is better suited on a particular switch. In this method, packets sent from the IP address A to IP address B, from IP address A to IP address C, and from IP address C to IP address B could all use different ports in the channel.

Different load-balancing methods have different advantages, and the choice of a particular load-balancing method should be based on the position of the switch in the network and the kind of traffic that needs to be load-distributed. In Figure 40-3, an EtherChannel from a switch that is aggregating data from four workstations communicates with a router. Because the router is a single-MAC-address device, source-based forwarding on the switch EtherChannel ensures that the switch uses all available bandwidth to the router. The router is configured for destination-based forwarding because the large number of workstations ensures that the traffic is evenly distributed from the router EtherChannel.

Use the option that provides the greatest variety in your configuration. For example, if the traffic on a channel is only going to a single MAC address, using the destination-MAC address always chooses the same link in the channel. Using source addresses or IP addresses might result in better load balancing.

Figure 40-3 Load Distribution and Forwarding Methods

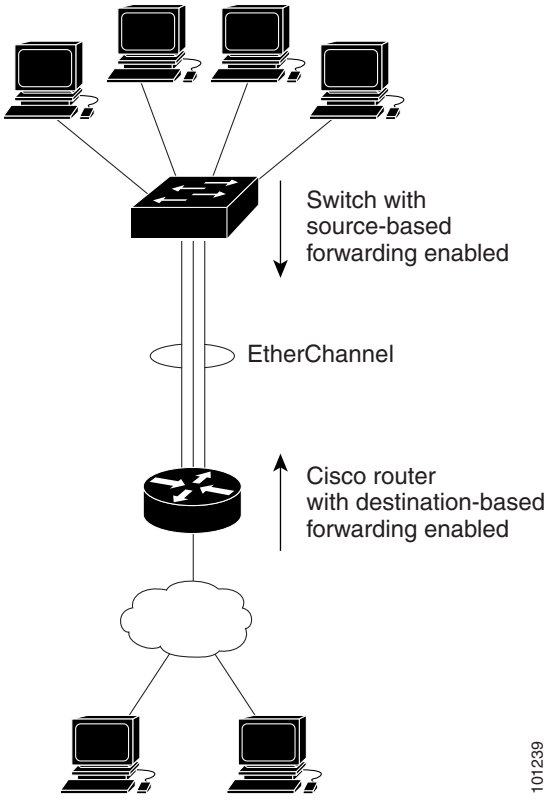

Default EtherChannel Settings

Table 40-3 Default EtherChannel Settings

Feature	Default Setting
Channel groups	None assigned.
Port-channel logical interface	None defined.
PAgP mode	No default.
PAgP learn method	Aggregate-port learning on all ports.
PAgP priority	128 on all ports.
LACP mode	No default.
LACP learn method	Aggregate-port learning on all ports.
LACP port priority	32768 on all ports.
LACP system priority	32768.
LACP system ID	LACP system priority and the switch MAC address.
Load balancing	Load distribution on the switch is based on the source-MAC address of the incoming packet.

EtherChannel Configuration Guidelines

If improperly configured, some EtherChannel ports are automatically disabled to avoid network loops and other problems. Follow these guidelines to avoid configuration problems:

- Do not try to configure more than 6 EtherChannels on the switch.
- Configure a PAgP EtherChannel with up to eight Ethernet ports of the same type.
- Configure a LACP EtherChannel with up to16 Ethernet ports of the same type. Up to eight ports can be active, and up to eight ports can be in standby mode.
- Configure all ports in an EtherChannel to operate at the same speeds and duplex modes.
- Enable all ports in an EtherChannel. A port in an EtherChannel that is disabled by using the **shutdown** interface configuration command is treated as a link failure, and its traffic is transferred to one of the remaining ports in the EtherChannel.
- When a group is first created, all ports follow the parameters set for the first port to be added to the group. If you change the configuration of one of these parameters, you must also make the changes to all ports in the group:
  - Allowed-VLAN list
  - Spanning-tree path cost for each VLAN
  - Spanning-tree port priority for each VLAN
  - Spanning-tree Port Fast setting
- Do not configure a port to be a member of more than one EtherChannel group.
- Do not configure an EtherChannel in both the PAgP and LACP modes. EtherChannel groups running PAgP and LACP can coexist on the same switch. Individual EtherChannel groups can run either PAgP or LACP, but they cannot interoperate.

- Do not configure a Switched Port Analyzer (SPAN) destination port as part of an EtherChannel.
- Do not configure a secure port as part of an EtherChannel or the reverse.
- Do not configure a private-VLAN port as part of an EtherChannel.
- Do not configure a port that is an active or a not-yet-active member of an EtherChannel as an IEEE 802.1x port. If you try to enable IEEE 802.1x on an EtherChannel port, an error message appears, and IEEE 802.1x is not enabled.
- If EtherChannels are configured on switch interfaces, remove the EtherChannel configuration from the interfaces before globally enabling IEEE 802.1x on a switch by using the **dot1x system-auth-control** global configuration command.
- For Layer 2 EtherChannels:
  - Assign all ports in the EtherChannel to the same VLAN, or configure them as trunks. Ports with different native VLANs cannot form an EtherChannel.
  - If you configure an EtherChannel from trunk ports, verify that the trunking mode (ISL or IEEE 802.1Q) is the same on all the trunks. Inconsistent trunk modes on EtherChannel ports can have unexpected results.
  - An EtherChannel supports the same allowed range of VLANs on all the ports in a trunking Layer 2 EtherChannel. If the allowed range of VLANs is not the same, the ports do not form an EtherChannel even when PAgP is set to the **auto** or **desirable** mode.
  - Ports with different spanning-tree path costs can form an EtherChannel if they are otherwise compatibly configured. Setting different spanning-tree path costs does not, by itself, make ports incompatible for the formation of an EtherChannel.

How to Configure EtherChannels

Note After you configure an EtherChannel, configuration changes applied to the port-channel interface apply to all the physical ports assigned to the port-channel interface, and configuration changes applied to the physical port affect only the port where you apply the configuration.

Configuring Layer 2 EtherChannels

You configure Layer 2 EtherChannels by assigning ports to a channel group with the **channel-group** interface configuration command. This command automatically creates the port-channel logical interface.

This required task explains how to configure a Layer 2 Ethernet port to a Layer 2 EtherChannel.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies a physical port, and enter interface configuration mode. Valid interfaces include physical ports. For a PAgP EtherChannel, you can configure up to eight ports of the same type and speed for the same group. For a LACP EtherChannel, you can configure up to 16 Ethernet ports of the same type. Up to eight ports can be active, and up to eight ports can be in standby mode.
Step 3	**switchport mode** {**access** | **trunk**} **switchport access vlan** *vlan-id*	Assigns all ports as static-access ports in the same VLAN, or configures them as trunks. If you configure the port as a static-access port, assign it to only one VLAN. The range is 1 to 4096.

	Command	Purpose
Step 4	**channel-group** *channel-group-number* **mode** {**auto** [**non-silent**] | **desirable** [**non-silent**] | **on**} | {**active** | **passive**}	Assigns the port to a channel group, and specifies the PAgP or the LACP mode. For *channel-group-number*, the range is 1 to 6. For **mode**, select one of these keywords: • **auto**—Enables PAgP only if a PAgP device is detected. It places the port into a passive negotiating state, in which the port responds to PAgP packets it receives but does not start PAgP packet negotiation. • **desirable**—Unconditionally enables PAgP. It places the port into an active negotiating state, in which the port starts negotiations with other ports by sending PAgP packets. • **on**—Forces the port to channel without PAgP or LACP. In the **on** mode, an EtherChannel exists only when a port group in the **on** mode is connected to another port group in the **on** mode. • **non-silent**—(Optional) If your switch is connected to a partner that is PAgP-capable, configure the switch port for nonsilent operation when the port is in the **auto** or **desirable** mode. If you do not Specifies **non-silent**, silent is assumed. The silent setting is for connections to file servers or packet analyzers. This setting allows PAgP to operate, to attach the port to a channel group, and to use the port for transmission. • **active**—Enables LACP only if a LACP device is detected. It places the port into an active negotiating state in which the port starts negotiations with other ports by sending LACP packets. • **passive**—Enables LACP on the port and places it into a passive negotiating state in which the port responds to LACP packets that it receives, but does not start LACP packet negotiation. For information on compatible modes for the switch and its partner, see the "PAgP Modes" section on page 40-4 and the "LACP Modes" section on page 40-6.
Step 5	**end**	Returns to privileged EXEC mode.

Configuring EtherChannel Load Balancing

This task is optional.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**port-channel load-balance** {**dst-ip** | **dst-mac** | **src-dst-ip** | **src-dst-mac** | **src-ip** | **src-mac**}	Configures an EtherChannel load-balancing method. The default is **src-mac**. Select one of these load-distribution methods: • **dst-ip**—Specifies the destination-host IP address. • **dst-mac**—Specifies the destination-host MAC address of the incoming packet. • **src-dst-ip**— Specifies the source-and-destination host-IP address. • **src-dst-mac**—Specifies the source-and-destination host-MAC address. • **src-ip**— Specifies the source-host IP address. • **src-mac**—Specifies the source-MAC address of the incoming packet.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring the PAgP Learn Method and Priority

This task is optional.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Specifies the port for transmission, and enter interface configuration mode.
Step 3	**pagp learn-method** physical-port	Selects the PAgP learning method. By default, **aggregation-port learning** is selected, which means the switch sends packets to the source by using any of the ports in the EtherChannel. With aggregate-port learning, it is not important on which physical port the packet arrives. Select **physical-port** to connect with another switch that is a physical learner. Make sure to configure the **port-channel load-balance** global configuration command to **src-mac** as described in the "Configuring EtherChannel Load Balancing" section on page 40-14. The learning method must be configured the same at both ends of the link.

	Command	Purpose
Step 4	**pagp port-priority** *priority*	Assigns a priority so that the selected port is chosen for packet transmission. For *priority*, the range is 0 to 255. The default is 128. The higher the priority, the more likely that the port will be used for PAgP transmission.
Step 5	**end**	Returns to privileged EXEC mode.

Configuring the LACP Hot-Standby Ports

This task is optional.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**lacp system-priority** *priority*	Configures the LACP system priority. For *priority*, the range is 1 to 65535. The default is 32768. The lower the value, the higher the system priority.
Step 3	**interface** *interface-id*	Specifies the port to be configured, and enters interface configuration mode.
Step 4	**lacp port-priority** ***priority***	Configures the LACP port priority. For *priority*, the range is 1 to 65535. The default is 32768. The lower the value, the more likely that the port will be used for LACP transmission.
Step 5	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining EtherChannels on the IE 2000 Switch

Command	Purpose
show etherchannel [*channel-group-number* {**detail** | **port** | **port-channel** | **protocol** | **summary**}] {**detail** | **load-balance** | **port** | **port-channel** | **protocol** | **summary**}	Displays EtherChannel information in a brief, detailed, and one-line summary form. Also displays the load-balance or frame-distribution scheme, port, port-channel, and protocol information.
show pagp [*channel-group-number*] {**counters** | **internal** | **neighbor**}	Displays PAgP information such as traffic information, the internal PAgP configuration, and neighbor information.
show pagp [*channel-group-number*] **dual-active**	Displays the dual-active detection status.
show lacp [*channel-group-number*] {**counters** | **internal** | **neighbor**}	Displays LACP information such as traffic information, the internal LACP configuration, and neighbor information.

Configuration Examples for Configuring EtherChannels

Configuring EtherChannels: Examples

This example shows how to configure an EtherChannel and assign two ports as static-access ports in VLAN 10 to channel 5 with the PAgP mode **desirable**:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet1/1 -2
Switch(config-if-range)# switchport mode access
Switch(config-if-range)# switchport access vlan 10
Switch(config-if-range)# channel-group 5 mode desirable non-silent
Switch(config-if-range)# end
```

This example shows how to configure an EtherChannel and assign two ports as static-access ports in VLAN 10 to channel 5 with the LACP mode **active**:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet1/1 -2
Switch(config-if-range)# switchport mode access
Switch(config-if-range)# switchport access vlan 10
Switch(config-if-range)# channel-group 5 mode active
Switch(config-if-range)# end
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 41

Configuring Static IP Unicast Routing

This chapter describes how to configure IP Version 4 (IPv4) static IP unicast routing on the switch. Static routing is supported only on switched virtual interfaces (SVIs) and not on physical interfaces. The switch does not support routing protocols.

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Restrictions for Static IP Unicast Routing

- By default, static IP routing is disabled on the switch unless the SDM template is modified to support static routing.
- To use this feature, the switch must be running the LAN Base image.

Information About Configuring Static IP Unicast Routing

Note When configuring routing parameters on the switch and to allocate system resources to maximize the number of unicast routes allowed, use the **sdm prefer lanbase-routing** global configuration command to set the Switch Database Management (SDM) feature to the routing template. For more information on the SDM templates, see Chapter 11, "Configuring SDM Templates" or see the **sdm prefer** command in the command reference for this release.

IP Routing

In some network environments, VLANs are associated with individual networks or subnetworks. In an IP network, each subnetwork is mapped to an individual VLAN. Configuring VLANs helps control the size of the broadcast domain and keeps local traffic local. However, network devices in different VLANs cannot communicate with one another without a Layer 3 device to route traffic between the VLANs, referred to as inter-VLAN routing. You configure one or more routers to route traffic to the appropriate destination VLAN.

Figure 41-1 shows a basic routing topology. Switch A is in VLAN 10, and Switch B is in VLAN 20. The router has an interface in each VLAN.

Figure 41-1 Routing Topology Example

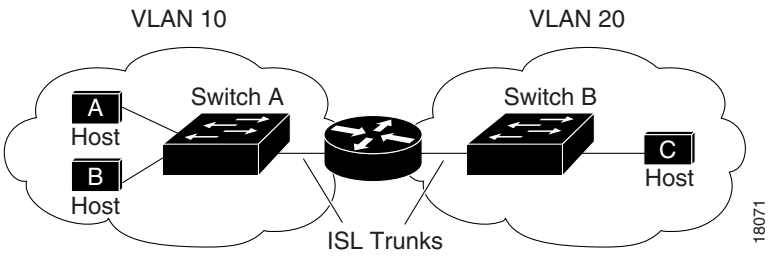

When Host A in VLAN 10 needs to communicate with Host B in VLAN 10, it sends a packet addressed to that host. Switch A forwards the packet directly to Host B, without sending it to the router.

When Host A sends a packet to Host C in VLAN 20, Switch A forwards the packet to the router, which receives the traffic on the VLAN 10 interface. The router uses the routing table to finds the correct outgoing interface, and forwards the packet on the VLAN 20 interface to Switch B. Switch B receives the packet and forwards it to Host C.

When static routing is enabled on Switch A and B, the router device is no longer needed to route packets.

Types of Routing

Routers and Layer 3 switches can route packets in these ways:

- Using default routing to send traffic with a destination unknown to the router to a default outlet or destination
- Using static routes to forward packets from predetermined ports through a single path into and out of a network
- Dynamically calculating routes by using a routing protocol

The switch supports static routes and default routes. It does not support routing protocols.

How to Configure Static IP Unicast Routing

Steps for Configuring Routing

In these procedures, the specified interface must be a switch virtual interface (SVI)—a VLAN interface created by using the **interface vlan** *vlan_id* global configuration command and by default a Layer 3 interface. All Layer 3 interfaces on which routing will occur must have IP addresses assigned to them. See the "Assigning IP Addresses to SVIs" section on page 41-3.

Note The switch supports 16 static routes (including user-configured routes and the default route) and any directly connected routes and default routes for the management interface. The switch can have an IP address assigned to each SVI. Before enabling routing, enter the **sdm prefer lanbase-routing** global configuration command and reload the switch.

Procedures for configuring routing:

- To support VLAN interfaces, create and configure VLANs on the switch, and assign VLAN membership to Layer 2 interfaces. For more information, see Chapter 17, "Configuring VLANs."
- Configure Layer 3 interfaces (SVIs) and physical routed port (no switchport).
- Assign IP addresses to the Layer 3 interfaces.
- Configure static routes

Enabling IP Unicast Routing

By default, the switch is in Layer 2 switching mode, and IP routing is disabled. To use the Layer 3 capabilities of the switch, enable IP routing.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip routing**	Enables IP routing.
Step 3	**end**	Returns to privileged EXEC mode.

Assigning IP Addresses to SVIs

To configure IP routing, you need to assign IP addresses to Layer 3 network interfaces. This enables communication with the hosts on those interfaces that use IP. IP routing is disabled by default, and no IP addresses are assigned to SVIs.

An IP address identifies a destination for IP packets. Some IP addresses are reserved for special uses and cannot be used for host, subnet, or network addresses. RFC 1166, "Internet Numbers," contains the official description of these IP addresses.

An interface can have one primary IP address. A a subnet mask identifies the bits that denote the network number in an IP address.

This task explains how to assign an IP address and a network mask to an SVI

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface vlan** *vlan_id*	Enters interface configuration mode, and specifies the Layer 3 VLAN to configure.
Step 3	**ip address** *ip-address subnet-mask*	Configures the IP address and IP subnet mask.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring Static Unicast Routes

Static unicast routes are user-defined routes that cause packets moving between a source and a destination to take a specified path. Static routes can be important if the router cannot build a route to a particular destination and are useful for specifying a gateway of last resort to which all unroutable packets are sent.

Use the **no ip route** *prefix mask* {*address* | *interface*} global configuration command to remove a static route. The switch retains static routes until you remove them.

When an interface goes down, all static routes through that interface are removed from the IP routing table. When the software can no longer find a valid next hop for the address specified as the forwarding router's address in a static route, the static route is also removed from the IP routing table.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip route** *prefix mask* {*address*	*interface*} [*distance*]	Establishs a static route.
Step 3	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining the IP Network

Command	Description
show interfaces [*interface-id*]	Displays the administrative and operational status of all interfaces specified interface.

Additional References for Configuring IP Unicast Routing

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
Cisco IOS IP address commands	*Cisco IOS IP Command Reference, Volume 1 of 3: Addressing and Services, Release 15.0*
Cisco IP routing configuration	*Cisco IOS IP Routing Configuration Guides, Release 15.0*
SDM template configuration	Chapter 11, "Configuring SDM Templates"

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 42

Configuring IPv6 Host Functions

Note To use IPv6 host functions, the switch must be running the LAN Base image.

This chapter describes how to configure IPv6 host functions on the switch.

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites Configuring IPv6 Host Functions

- To enable dual-stack environments (supporting both IPv4 and IPv6), you must configure the switch to use the a dual IPv4 and IPv6 switch database management (SDM) template. See the "Dual IPv4 and IPv6 Protocol Stacks" section on page 42-4.

Information About Configuring IPv6 Host Functions

IPv6

IPv4 users can move to IPv6 and receive services such as end-to-end security, quality of service (QoS), and globally unique addresses. The IPv6 address space reduces the need for private addresses and Network Address Translation (NAT) processing by border routers at network edges.

For information about how Cisco Systems implements IPv6, go to this URL:

http://www.cisco.com/en/US/products/ps6553/products_ios_technology_home.html

For information about IPv6 and other features in this chapter

- See the *Cisco IOS IPv6 Configuration Library* at this URL:

  http://www.cisco.com/en/US//docs/ios-xml/ios/ipv6/configuration/15-1mt/ipv6-15-1mt-book.html

This section describes IPv6 implementation on the switch. These sections are included:

- IPv6 Addresses, page 42-2
- Supported IPv6 Host Features, page 42-2
- How to Configure IPv6 Hosting, page 42-7

IPv6 Addresses

The switch supports only IPv6 unicast addresses. It does not support site-local unicast addresses, anycast addresses, or multicast addresses.

The IPv6 128-bit addresses are represented as a series of eight 16-bit hexadecimal fields separated by colons in the format: n:n:n:n:n:n:n:n. This is an example of an IPv6 address:

2031:0000:130F:0000:0000:09C0:080F:130B

For easier implementation, leading zeros in each field are optional. This is the same address without leading zeros:

2031:0:130F:0:0:9C0:80F:130B

You can also use two colons (::) to represent successive hexadecimal fields of zeros, but you can use this short version only once in each address:

2031:0:130F::09C0:080F:130B

For more information about IPv6 address formats, address types, and the IPv6 packet header, see the "Implementing IPv6 Addressing and Basic Connectivity" chapter of *Cisco IOS IPv6 Configuration Library* on Cisco.com.

In the "Implementing Addressing and Basic Connectivity" chapter, these sections apply to the switch:

- IPv6 Address Formats
- IPv6 Address Output Display
- Simplified IPv6 Packet Header

Supported IPv6 Host Features

These sections describe the IPv6 protocol features supported by the switch:

- 128-Bit Wide Unicast Addresses, page 42-3
- DNS for IPv6, page 42-3
- ICMPv6, page 42-3
- Neighbor Discovery, page 42-3
- Default Router Preference, page 42-4
- IPv6 Stateless Autoconfiguration and Duplicate Address Detection, page 42-4
- IPv6 Applications, page 42-4
- Dual IPv4 and IPv6 Protocol Stacks, page 42-4

- SNMP and Syslog Over IPv6, page 42-5
- HTTP over IPv6, page 42-6

Support on the switch includes expanded address capability, header format simplification, improved support of extensions and options, and hardware parsing of the extension header. The switch supports hop-by-hop extension header packets, which are routed or bridged in software.

128-Bit Wide Unicast Addresses

The switch supports aggregatable global unicast addresses and link-local unicast addresses. It does not support site-local unicast addresses.

- Aggregatable global unicast addresses are IPv6 addresses from the aggregatable global unicast prefix. The address structure enables strict aggregation of routing prefixes and limits the number of routing table entries in the global routing table. These addresses are used on links that are aggregated through organizations and eventually to the Internet service provider.

  These addresses are defined by a global routing prefix, a subnet ID, and an interface ID. Current global unicast address allocation uses the range of addresses that start with binary value 001 (2000::/3). Addresses with a prefix of 2000::/3(001) through E000::/3(111) must have 64-bit interface identifiers in the extended unique identifier (EUI)-64 format.

- Link local unicast addresses can be automatically configured on any interface by using the link-local prefix FE80::/10(1111 1110 10) and the interface identifier in the modified EUI format. Link-local addresses are used in the neighbor discovery protocol (NDP) and the stateless autoconfiguration process. Nodes on a local link use link-local addresses and do not require globally unique addresses to communicate. IPv6 routers do not forward packets with link-local source or destination addresses to other links.

For more information, see the section about IPv6 unicast addresses in the “Implementing IPv6 Addressing and Basic Connectivity” chapter in the *Cisco IOS IPv6 Configuration Library* on Cisco.com.

DNS for IPv6

IPv6 supports Domain Name System (DNS) record types in the DNS name-to-address and address-to-name lookup processes. The DNS AAAA resource record types support IPv6 addresses and are equivalent to an A address record in IPv4. The switch supports DNS resolution for IPv4 and IPv6.

ICMPv6

The Internet Control Message Protocol (ICMP) in IPv6 generates error messages, such as ICMP destination unreachable messages, to report errors during processing and other diagnostic functions. In IPv6, ICMP packets are also used in the neighbor discovery protocol and path MTU discovery.

Neighbor Discovery

The switch supports NDP for IPv6, a protocol running on top of ICMPv6, and static neighbor entries for IPv6 stations that do not support NDP. The IPv6 neighbor discovery process uses ICMP messages and solicited-node multicast addresses to determine the link-layer address of a neighbor on the same network (local link), to verify the reachability of the neighbor, and to keep track of neighboring routers.

The switch supports ICMPv6 redirect for routes with mask lengths less than 64 bits. ICMP redirect is not supported for host routes or for summarized routes with mask lengths greater than 64 bits.

Neighbor discovery throttling ensures that the switch CPU is not unnecessarily burdened while it is in the process of obtaining the next hop forwarding information to route an IPv6 packet. The switch drops any additional IPv6 packets whose next hop is the same neighbor that the switch is actively trying to resolve. This drop avoids further load on the CPU.

Default Router Preference

The switch supports IPv6 default router preference (DRP), an extension in router advertisement messages. DRP improves the ability of a host to select an appropriate router, especially when the host is multihomed and the routers are on different links. The switch does not support the Route Information Option in RFC 4191.

An IPv6 host maintains a default router list from which it selects a router for traffic to offlink destinations. The selected router for a destination is then cached in the destination cache. NDP for IPv6 specifies that routers that are reachable or probably reachable are preferred over routers whose reachability is unknown or suspect. For reachable or probably reachable routers, NDP can either select the same router every time or cycle through the router list. By using DRP, you can configure an IPv6 host to prefer one router over another, provided both are reachable or probably reachable.

For more information about DRP for IPv6, see the "Implementing IPv6 Addresses and Basic Connectivity" chapter in the *Cisco IOS IPv6 Configuration Library* on Cisco.com.

IPv6 Stateless Autoconfiguration and Duplicate Address Detection

The switch uses stateless autoconfiguration to manage link, subnet, and site addressing changes, such as management of host and mobile IP addresses. A host autonomously configures its own link-local address, and booting nodes send router solicitations to request router advertisements for configuring interfaces.

For more information about autoconfiguration and duplicate address detection, see the "Implementing IPv6 Addressing and Basic Connectivity" chapter of *Cisco IOS IPv6 Configuration Library* on Cisco.com.

IPv6 Applications

The switch has IPv6 support for these applications:

- Ping, traceroute, Telnet, TFTP, and FTP
- Secure Shell (SSH) over an IPv6 transport
- HTTP server access over IPv6 transport
- DNS resolver for AAAA over IPv4 transport
- Cisco Discovery Protocol (CDP) support for IPv6 addresses

For more information about managing these applications, see the "Managing Cisco IOS Applications over IPv6" chapter and the "Implementing IPv6 Addressing and Basic Connectivity" chapter in the *Cisco IOS IPv6 Configuration Library* on Cisco.com.

Dual IPv4 and IPv6 Protocol Stacks

You must use the dual IPv4 and IPv6 template to allocate ternary content addressable memory (TCAM) usage to both IPv4 and IPv6 protocols.

Figure 42-1 shows a router forwarding both IPv4 and IPv6 traffic through the same interface, based on the IP packet and destination addresses.

Figure 42-1 Dual IPv4 and IPv6 Support on an Interface

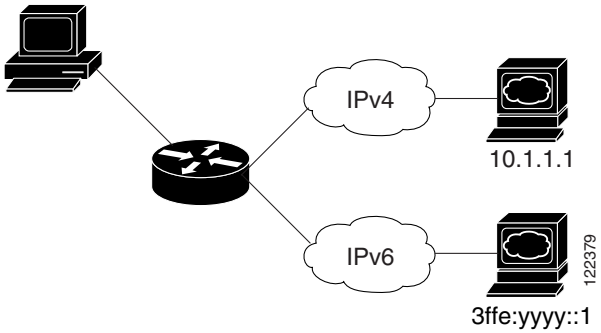

Use the dual IPv4 and IPv6 switch database management (SDM) template to enable dual-stack environments (supporting both IPv4 and IPv6). For more information about the dual IPv4 and IPv6 SDM template, see Chapter 11, "Configuring SDM Templates."

The dual IPv4 and IPv6 templates allow the switch to be used in dual-stack environments.

- If you try to configure IPv6 without first selecting a dual IPv4 and IPv6 template, a warning message appears.
- In IPv4-only environments, the switch applies IPv4 QoS and ACLs in hardware. IPv6 packets are not supported.
- In dual IPv4 and IPv6 environments, the switch applies IPv4 QoS and ACLs in hardware.
- IPv6 QoS and ACLs are not supported.
- If you do not plan to use IPv6, do not use the dual-stack template because this template results in less TCAM capacity for each resource.

For more information about IPv4 and IPv6 protocol stacks, see the "Implementing IPv6 Addressing and Basic Connectivity" chapter of *Cisco IOS IPv6 Configuration Library* on Cisco.com.

Static Routes for IPv6

Static routes are manually configured and define an explicit route between two networking devices. Static routes are useful for smaller networks with only one path to an outside network or to provide security for certain types of traffic in a larger network.

For more information about static routes, see the "Implementing Static Routes for IPv6" chapter in the *Cisco IOS IPv6 Configuration Library* on Cisco.com.

SNMP and Syslog Over IPv6

To support both IPv4 and IPv6, IPv6 network management requires both IPv6 and IPv4 transports. Syslog over IPv6 supports address data types for these transports.

SNMP and syslog over IPv6 provide these features:

- Support for both IPv4 and IPv6
- IPv6 transport for SNMP and to modify the SNMP agent to support traps for an IPv6 host
- SNMP- and syslog-related MIBs to support IPv6 addressing

- Configuration of IPv6 hosts as trap receivers

For support over IPv6, SNMP modifies the existing IP transport mapping to simultaneously support IPv4 and IPv6. These SNMP actions support IPv6 transport management:

- Opens User Datagram Protocol (UDP) SNMP socket with default settings
- Provides a new transport mechanism called *SR_IPV6_TRANSPORT*
- Sends SNMP notifications over IPv6 transport
- Supports SNMP-named access lists for IPv6 transport
- Supports SNMP proxy forwarding using IPv6 transport
- Verifies SNMP Manager feature works with IPv6 transport

For information on SNMP over IPv6, including configuration procedures, see the "Managing Cisco IOS Applications over IPv6" chapter in the *Cisco IOS IPv6 Configuration Library* on Cisco.com.

For information about syslog over IPv6, including configuration procedures, see the "Implementing IPv6 Addressing and Basic Connectivity" chapter in the *Cisco IOS IPv6 Configuration Library* on Cisco.com.

HTTP over IPv6

The HTTP client sends requests to both IPv4 and IPv6 HTTP servers, which respond to requests from both IPv4 and IPv6 HTTP clients. URLs with literal IPv6 addresses must be specified in hexadecimal using 16-bit values between colons.

The accept socket call chooses an IPv4 or IPv6 address family. The accept socket is either an IPv4 or IPv6 socket. The listening socket continues to listen for both IPv4 and IPv6 signals that indicate a connection. The IPv6 listening socket is bound to an IPv6 wildcard address.

The underlying TCP/IP stack supports a dual-stack environment. HTTP relies on the TCP/IP stack and the sockets for processing network-layer interactions.

Basic network connectivity (**ping**) must exist between the client and the server hosts before HTTP connections can be made.

Default IPv6 Settings

Table 42-1 Default IPv6 Settings

Feature	Default Setting
SDM template	Default.
IPv6 addresses	None configured.

How to Configure IPv6 Hosting

Configuring IPv6 Addressing and Enabling IPv6 Host

This section describes how to assign IPv6 addresses to individual Layer 3 interfaces and to globally forward IPv6 traffic on the switch.

Before configuring IPv6 on the switch, consider these guidelines:

- Be sure to select a dual IPv4 and IPv6 SDM template.
- In the **ipv6 address** interface configuration command, you must enter the *ipv6-address* and *ipv6-prefix* variables with the address specified in hexadecimal using 16-bit values between colons. The *prefix-length* variable (preceded by a slash [/]) is a decimal value that shows how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address).

To forward IPv6 traffic on an interface, you must configure a global IPv6 address on that interface. Configuring an IPv6 address on an interface automatically configures a link-local address and activates IPv6 for the interface. The configured interface automatically joins these required multicast groups for that link:

- solicited-node multicast group FF02:0:0:0:0:1:ff00::/104 for each unicast address assigned to the interface (this address is used in the neighbor discovery process.)
- all-nodes link-local multicast group FF02::1
- all-routers link-local multicast group FF02::2

For more information about configuring IPv6, see the “Implementing Addressing and Basic Connectivity for IPv6” chapter in the *Cisco IOS IPv6 Configuration Library* on Cisco.com.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**sdm prefer dual-ipv4-and-ipv6 default**	Selects the SDM template that supports IPv4 and IPv6.
Step 3	**end**	Returns to privileged EXEC mode.
Step 4	**reload**	Reloads the operating system.
Step 5	**configure terminal**	Enters global configuration mode after the switch reloads.
Step 6	**interface** *interface-id*	Enters interface configuration mode, and specifies the interface to configure.

	Command	Purpose
Step 7	**ipv6 address** *ipv6-prefix/prefix length* **eui-64** or **ipv6 address** *ipv6-address* **link-local** or **ipv6 enable**	• Specifies a global IPv6 address with an extended unique identifier (EUI) in the low-order 64 bits of the IPv6 address. • Specifies only the network prefix; the last 64 bits are automatically computed from the switch MAC address. This enables IPv6 processing on the interface. • Specifies a link-local address on the interface to be used instead of the link-local address that is automatically configured when IPv6 is enabled on the interface. This command enables IPv6 processing on the interface. • Automatically configures an IPv6 link-local address on the interface, and enable the interface for IPv6 processing. The link-local address can only be used to communicate with nodes on the same link.
Step 8	**exit**	Returns to global configuration mode.
Step 9	**end**	Returns to privileged EXEC mode.

Configuring Default Router Preference

Router advertisement messages are sent with the default router preference (DRP) configured by the **ipv6 nd router-preference** interface configuration command. If no DRP is configured, RAs are sent with a medium preference.

A DRP is useful when two routers on a link might provide equivalent, but not equal-cost routing, and policy might dictate that hosts should prefer one of the routers.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface** *interface-id*	Enters interface configuration mode, and enters the Layer 3 interface on which you want to specify the DRP.
Step 3	**ipv6 nd router-preference** {**high** | **medium** | **low**}	Specifies a DRP for the router on the switch interface.
Step 4	**end**	Returns to privileged EXEC mode.

Configuring IPv6 ICMP Rate Limiting

ICMP rate limiting is enabled by default with a default interval between error messages of 100 milliseconds and a bucket size (maximum number of tokens to be stored in a bucket) of 10.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ipv6 icmp error-interval** *interval* [*bucketsize*]	Configures the interval and bucket size for IPv6 ICMP error messages: • *interval*—The interval (in milliseconds) between tokens being added to the bucket. The range is from 0 to 2147483647 milliseconds. • *bucketsize*—(Optional) The maximum number of tokens stored in the bucket. The range is from 1 to 200.
Step 3	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining IPv6 Host Information

Command	Purpose
show ipv6 interface *interface-id*	Displays IPv6 interface status and configuration.
show ipv6 mtu	Displays IPv6 MTU per destination cache.
show ipv6 neighbors	Displays IPv6 neighbor cache entries.
show ipv6 prefix-list	Displays a list of IPv6 prefix lists.
show ipv6 protocols	Displays IPv6 routing protocols on the switch.
show ipv6 route	Displays the IPv6 route table entries.
show ipv6 static	Displays IPv6 static routes.
show ipv6 traffic	Displays IPv6 traffic statistics.
show ip http server history	Displays the previous 20 connections to the HTTP server, including the IP address accessed and the time when the connection was closed.
show ip http server connection	Displays the current connections to the HTTP server, including the local and remote IP addresses being accessed.
show ip http client connection	Displays the configuration values for HTTP client connections to HTTP servers.
show ip http client history	Displays a list of the last 20 requests made by the HTTP client to the server.

Configuration Examples for IPv6 Host Functions

Enabling IPv6: Example

This example shows how to enable IPv6 with both a link-local address and a global address based on the IPv6 prefix 2001:0DB8:c18:1::/64. The EUI-64 interface ID is used in the low-order 64 bits of both addresses. Output from the **show ipv6 interface** EXEC command shows how the interface ID (20B:46FF:FE2F:D940) is appended to the link-local prefix FE80::/64 of the interface.

```
Switch(config)# sdm prefer dual-ipv4-and-ipv6 default
Switch(config)# interface gigabitethernetfastethernet1/0/11
Switch(config-if)# ipv6 address 2001:0DB8:c18:1::/64 eui 64
Switch(config-if)# end
Switch# show ipv6 interface gigabitethernetfastethernet1/0/11
GigabitEthernetFastEthernet1/0/11 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::20B:46FF:FE2F:D940
  Global unicast address(es):
  2001:0DB8:c18:1:20B:46FF:FE2F:D940, subnet is 2001:0DB8:c18:1::/64 [EUI]
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF2F:D940
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds
  ND advertised reachable time is 0 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses.
```

Configuring DRP: Example

This example shows how to configure a DRP of *high* for the router on an interface.

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# ipv6 nd router-preference high
Switch(config-if)# end
```

Configuring an IPv6 ICMP Error Message Interval

This example shows how to configure an IPv6 ICMP error message interval of 50 milliseconds and a bucket size of 20 tokens.

```
Switch(config)# ipv6 icmp error-interval 50 20
```

Displaying Show Command Output: Examples

This is an example of the output from the **show ipv6 interface** privileged EXEC command:

```
Switch# show ipv6 interface
Vlan1 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::20B:46FF:FE2F:D940
  Global unicast address(es):
    3FFE:C000:0:1:20B:46FF:FE2F:D940, subnet is 3FFE:C000:0:1::/64 [EUI]
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF2F:D940
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds
  ND advertised reachable time is 0 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
<output truncated>
```

This is an example of the output from the **show ipv6 protocols** privileged EXEC command:

```
Switch# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "static"
IPv6 Routing Protocol is "rip fer"
  Interfaces:
    Vlan6
    FastEthernet0/4
    FastEthernet0/11
    FastEthernet0/12
    GigabitEthernet2/0/4
    GigabitEthernet2/0/
    GigabitEthernet1/0/12
Redistribution:
    None
```

This is an example of the output from the **show ipv6 neighbor** privileged EXEC command:

```
Switch# show ipv6 neighbors
IPv6 Address                              Age Link-layer Addr State Interface
3FFE:C000:0:7::777                          - 0007.0007.0007  REACH Vl7
3FFE:C101:113:1::33                         - 0000.0000.0033  REACH Fa1/0/13
```

This is an example of the output from the **show ipv6 route** privileged EXEC command:

```
Switch# show ipv6 route
IPv6 Routing Table - Default - 1 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
L   FF00::/8 [0/0]
     via Null0, receive
```

This is an example of the output from the **show ipv6 traffic** privileged EXEC command.

```
Switch# show ipv6 traffic
IPv6 statistics:
  Rcvd:  1 total, 1 local destination
         0 source-routed, 0 truncated
         0 format errors, 0 hop count exceeded
         0 bad header, 0 unknown option, 0 bad source
```

```
          0 unknown protocol, 0 not a router
          0 fragments, 0 total reassembled
          0 reassembly timeouts, 0 reassembly failures
  Sent:  36861 generated, 0 forwarded
          0 fragmented into 0 fragments, 0 failed
          0 encapsulation failed, 0 no route, 0 too big
          0 RPF drops, 0 RPF suppressed drops
  Mcast: 1 received, 36861 sent

ICMP statistics:
  Rcvd: 1 input, 0 checksum errors, 0 too short
        0 unknown info type, 0 unknown error type
        unreach: 0 routing, 0 admin, 0 neighbor, 0 address, 0 port
        parameter: 0 error, 0 header, 0 option
        0 hopcount expired, 0 reassembly timeout,0 too big
        0 echo request, 0 echo reply
        0 group query, 0 group report, 0 group reduce
        1 router solicit, 0 router advert, 0 redirects
        0 neighbor solicit, 0 neighbor advert
  Sent: 10112 output, 0 rate-limited
        unreach: 0 routing, 0 admin, 0 neighbor, 0 address, 0 port
        parameter: 0 error, 0 header, 0 option
        0 hopcount expired, 0 reassembly timeout,0 too big
        0 echo request, 0 echo reply
        0 group query, 0 group report, 0 group reduce
        0 router solicit, 9944 router advert, 0 redirects
        84 neighbor solicit, 84 neighbor advert

UDP statistics:
  Rcvd: 0 input, 0 checksum errors, 0 length errors
        0 no port, 0 dropped
  Sent: 26749 output

TCP statistics:
  Rcvd: 0 input, 0 checksum errors
  Sent: 0 output, 0 retransmitted
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
Cisco IOS static IPv6 routing	"Implementing Static Routes for IPv6" chapter in the *Cisco IOS IPv6 Configuration Library* on Cisco.com.
DRP for IPv6	"Implementing IPv6 Addresses and Basic Connectivity" chapter in the *Cisco IOS IPv6 Configuration Library* on Cisco.com

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 43

Configuring Link State Tracking

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Restrictions for Configuring Link State Tracking

- To use this feature, the switch must be running the LAN Base image.
- An interface that is defined as an upstream interface cannot also be defined as a downstream interface in the same or a different link state group. The reverse is also true.
- An interface cannot be a member of more than one link state group.
- You can configure only two link state groups per switch.

Information About Configuring Link State Tracking

Link State Tracking

Link state tracking, also known as trunk failover, is a feature that binds the link state of multiple interfaces. For example, link state tracking provides redundancy in the network when used with server NIC adapter teaming. When the server network adapters are configured in a primary or secondary relationship known as teaming, if the link is lost on the primary interface, connectivity is transparently changed to the secondary interface.

Note An interface can be an aggregation of ports (an EtherChannel), a single physical port in access or trunk mode, or a routed port.

Figure 43-1 on page 43-3 shows a network configured with link state tracking. To enable link state tracking, create a *link state group*, and specify the interfaces that are assigned to the link state group. In a link state group, these interfaces are bundled together. The *downstream interfaces* are bound to the *upstream interfaces*. Interfaces connected to servers are referred to as downstream interfaces, and interfaces connected to distribution switches and network devices are referred to as upstream interfaces.

The configuration in Figure 43-1 ensures that the network traffic flow is balanced as follows:

- For links to switches and other network devices
  - Server 1 and server 2 use switch A for primary links and switch B for secondary links.
  - Server 3 and server 4 use switch B for primary links and switch A for secondary links.
- Link state group 1 on switch A
  - Switch A provides primary links to server 1 and server 2 through link state group 1. Port 1 is connected to server 1, and port 2 is connected to server 2. Port 1 and port 2 are the downstream interfaces in link state group 1.
  - Port 5 and port 6 are connected to distribution switch 1 through link state group 1. Port 5 and port 6 are the upstream interfaces in link state group 1.
- Link state group 2 on switch A
  - Switch A provides secondary links to server 3 and server 4 through link state group 2. Port 3 is connected to server 3, and port 4 is connected to server 4. Port 3 and port 4 are the downstream interfaces in link state group 2.
  - Port 7 and port 8 are connected to distribution switch 2 through link state group 2. Port 7 and port 8 are the upstream interfaces in link state group 2.
- Link state group 2 on switch B
  - Switch B provides primary links to server 3 and server 4 through link state group 2. Port 3 is connected to server 3, and port 4 is connected to server 4. Port 3 and port 4 are the downstream interfaces in link state group 2.
  - Port 5 and port 6 are connected to distribution switch 2 through link state group 2. Port 5 and port 6 are the upstream interfaces in link state group 2.
- Link state group 1 on switch B
  - Switch B provides secondary links to server 1 and server 2 through link state group 1. Port 1 is connected to server 1, and port 2 is connected to server 2. Port 1 and port 2 are the downstream interfaces in link state group 1.
  - Port 7 and port 8 are connected to distribution switch 1 through link state group 1. Port 7 and port 8 are the upstream interfaces in link state group 1.

In a link state group, the upstream ports can become unavailable or lose connectivity because the distribution switch or router fails, the cables are disconnected, or the link is lost. These are the interactions between the downstream and upstream interfaces when link state tracking is enabled:

- If any of the upstream interfaces are in the link-up state, the downstream interfaces can change to or remain in the link-up state.
- If all of the upstream interfaces become unavailable, link state tracking automatically puts the downstream interfaces in the error-disabled state. Connectivity to and from the servers is automatically changed from the primary server interface to the secondary server interface.

  As an example of a connectivity change from link state group 1 to link state group 2 on switch A, see Figure 43-1 on page 43-3. If the upstream link for port 6 is lost, the link states of downstream ports 1 and 2 do not change. However, if the link for upstream port 5 is also lost, the link state of the

downstream ports changes to the link-down state. Connectivity to server 1 and server 2 is then changed from link state group1 to link state group 2. The downstream ports 3 and 4 do not change state because they are in link-group 2.

- If the link state group is configured, link state tracking is disabled, and the upstream interfaces lose connectivity, the link states of the downstream interfaces remain unchanged. The server does not recognize that upstream connectivity has been lost and does not failover to the secondary interface.

You can recover a downstream interface link-down condition by removing the failed downstream port from the link state group. To recover multiple downstream interfaces, disable the link state group.

Figure 43-1 Typical Link State Tracking Configuration

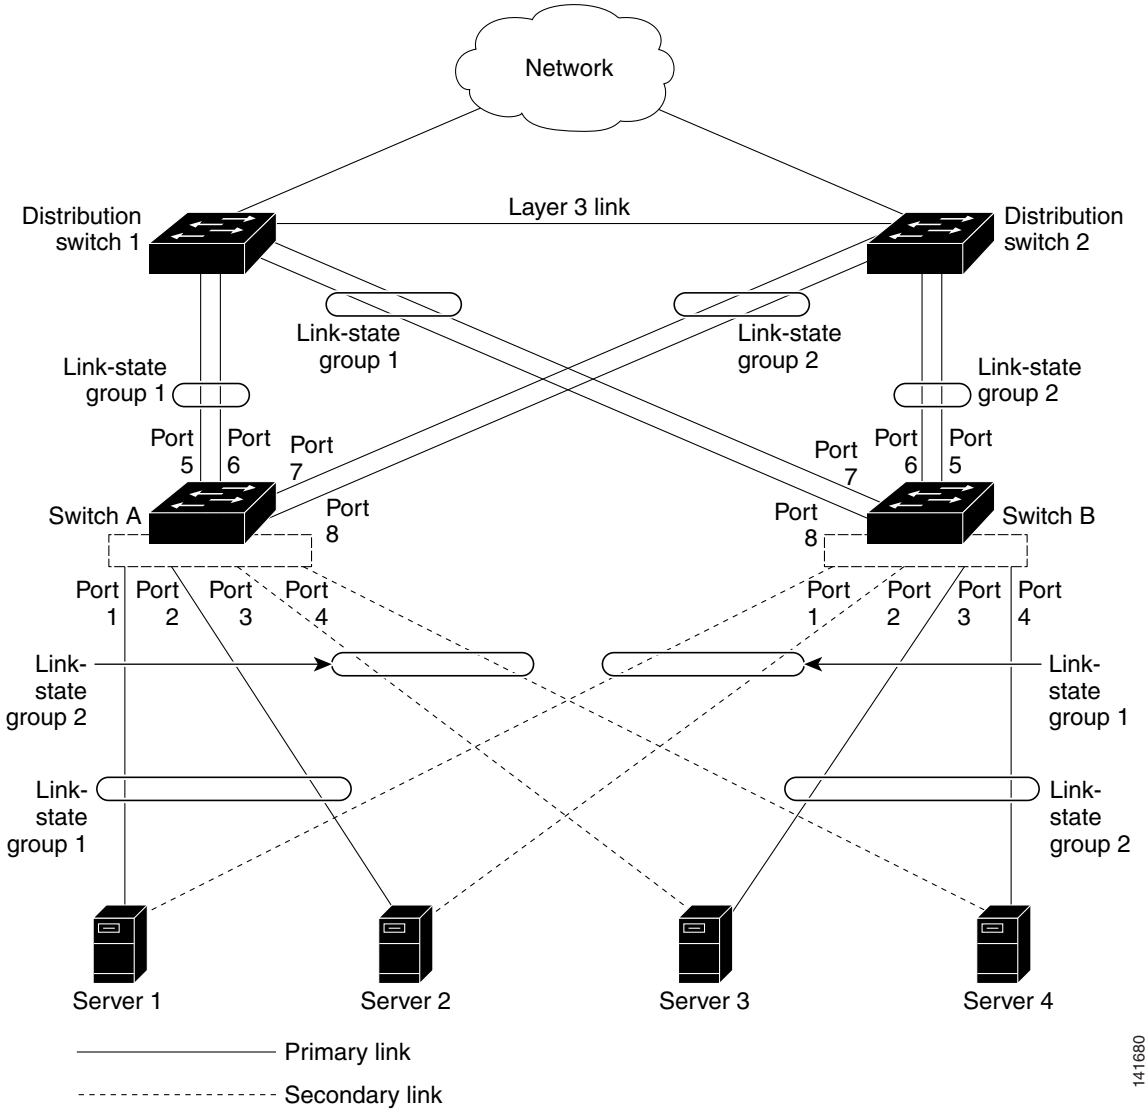

Default Link State Tracking Configuration

There are no link state groups defined, and link state tracking is not enabled for any group.

How to Configure Link State Tracking

Configuring Link State Tracking

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**link state track** *number*	Creates a link state group, and enables link state tracking. The group number can be 1 to 2; the default is 1.
Step 3	**interface** *interface-id*	Specifies a physical interface or range of interfaces to configure, and enters interface configuration mode. Valid interfaces include switch ports in access or trunk mode (IEEE 802.1q), routed ports, or multiple ports bundled into an EtherChannel interface (static or LACP), also in trunk mode.
Step 4	**link state group** [*number*] {**upstream** | **downstream**}	Specifies a link state group, and configures the interface as either an **upstream** or **downstream** interface in the group.The group number can be 1 to 2; the default is 1.
Step 5	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining Link State Tracking

Command	Purpose
show link state group	Displays the link state group information.

Configuration Examples for Configuring Link State Tracking

Displaying Link State Information: Examples

Use the **show link state group** command to display the link state group information. Enter this command without keywords to display information about all link state groups. Enter the group number to display information specific to the group. Enter the detail keyword to display detailed information about the group.

This is an example of output from the **show link state group 1** command:

```
Switch> show link state group 1

Link State Group: 1      Status: Enabled, Down
```

This is an example of output from the **show link state group detail** command:

```
Switch> show link state group detail

(Up):Interface up   (Dwn):Interface Down   (Dis):Interface disabled
```

```
Link State Group: 1 Status: Enabled, Down
Upstream Interfaces : Fa1/7(Dwn) Fa1/8(Dwn)
Downstream Interfaces : Fa1/3(Dis) Fa1/4(Dis) Fa1/5(Dis) Fa1/6(Dis)

Link State Group: 2 Status: Enabled, Down
Upstream Interfaces : Fa1/6(Dwn) Fa1/7(Dwn) Fa1/8(Dwn)
Downstream Interfaces : Fa1/2(Dis) Fa1/3(Dis) Fa1/4(Dis) Fa1/5(Dis)

(Up):Interface up (Dwn):Interface Down (Dis):Interface disabled
```

Creating a Link State Group: Example

This example shows how to create a link state group and configure the interfaces:

```
Switch# configure terminal
Switch(config)# link state track 1
Switch(config)# interface range gigabitethernet1/1 -2
Switch(config-if)# link state group 1 upstream
Switch(config-if)# interface gigabitethernet1/1
Switch(config-if)# link state group 1 downstream
Switch(config-if)# interface gigabitethernet1/1
Switch(config-if)# link state group 1 downstream
Switch(config-if)# interface gigabitethernet1/2
Switch(config-if)# link state group 1 downstream
Switch(config-if)# end
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
EtherChannel configuration	*Chapter 40, "Configuring EtherChannels"*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 44

Configuring IPv6 MLD Snooping

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for Configuring IPv6 MLD Snooping

- To use IPv6, you must configure the dual IPv4 and IPv6 Switch Database Management (SDM) template on the switch. You select the template by entering the **sdm prefer dual-ipv4-and-ipv6** global configuration command.

Restrictions for Configuring IPv6 MLD Snooping

- To use this feature, the switch must be running the LAN Base image.
- You can use Multicast Listener Discovery (MLD) snooping to enable efficient distribution of IP version 6 (IPv6) multicast data to clients and routers in a switched network on the switch.

Information About Configuring IPv6 MLD Snooping

IPv6 MLD Snooping

In IP version 4 (IPv4), Layer 2 switches can use Internet Group Management Protocol (IGMP) snooping to limit the flooding of multicast traffic by dynamically configuring Layer 2 interfaces so that multicast traffic is forwarded to only those interfaces associated with IP multicast devices. In IPv6, MLD snooping performs a similar function. With MLD snooping, IPv6 multicast data is selectively forwarded to a list of ports that want to receive the data, instead of being flooded to all ports in a VLAN. This list is constructed by snooping IPv6 multicast control packets.

MLD is a protocol used by IPv6 multicast routers to discover the presence of multicast listeners (nodes that want to receive IPv6 multicast packets) on its directly attached links and to discover which multicast packets are of interest to neighboring nodes. MLD is derived from IGMP; MLD version 1 (MLDv1) is equivalent to IGMPv2 and MLD version 2 (MLDv2) is equivalent to IGMPv3. MLD is a subprotocol of Internet Control Message Protocol version 6 (ICMPv6), and MLD messages are a subset of ICMPv6 messages, identified in IPv6 packets by a preceding Next Header value of 58.

The switch supports two versions of MLD snooping:

- MLDv1 snooping detects MLDv1 control packets and sets up traffic bridging based on IPv6 destination multicast addresses.
- MLDv2 basic snooping (MBSS) uses MLDv2 control packets to set up traffic forwarding based on IPv6 destination multicast addresses.

The switch can snoop on both MLDv1 and MLDv2 protocol packets and bridge IPv6 multicast data based on destination IPv6 multicast addresses.

Note The switch does not support MLDv2 enhanced snooping (MESS), which sets up IPv6 source and destination multicast address-based forwarding.

MLD snooping can be enabled or disabled globally or per VLAN. When MLD snooping is enabled, a per-VLAN IPv6 multicast MAC address table is constructed in software and a per-VLAN IPv6 multicast address table is constructed in software and hardware. The switch then performs IPv6 multicast-address based bridging in hardware.

MLD Messages

MLDv1 supports three types of messages:

- Listener Queries are the equivalent of IGMPv2 queries and are either General Queries or Multicast-Address-Specific Queries (MASQs).
- Multicast Listener Reports are the equivalent of IGMPv2 reports.
- Multicast Listener Done messages are the equivalent of IGMPv2 leave messages.

MLDv2 supports MLDv2 queries and reports, as well as MLDv1 Report and Done messages.

Message timers and state transitions resulting from messages being sent or received are the same as those of IGMPv2 messages. MLD messages that do not have valid link-local IPv6 source addresses are ignored by MLD routers and switches.

MLD Queries

The switch sends out MLD queries, constructs an IPv6 multicast address database, and generates MLD group-specific and MLD group-and-source-specific queries in response to MLD Done messages. The switch also supports report suppression, report proxying, Immediate-Leave functionality, and static IPv6 multicast MAC-address configuration.

When MLD snooping is disabled, all MLD queries are flooded in the ingress VLAN.

When MLD snooping is enabled, received MLD queries are flooded in the ingress VLAN, and a copy of the query is sent to the CPU for processing. From the received query, MLD snooping builds the IPv6 multicast address database. It detects multicast router ports, maintains timers, sets report response time, learns the querier IP source address for the VLAN, learns the querier port in the VLAN, and maintains multicast-address aging.

Note When the IPv6 multicast router is a Catalyst 6500 switch and you are using extended VLANs (in the range 1006 to 4096), IPv6 MLD snooping must be enabled on the extended VLAN on the Catalyst 6500 switch in order for the switch to receive queries on the VLAN. For normal-range VLANs (1 to 1005), it is not necessary to enable IPv6 MLD snooping on the VLAN on the Catalyst 6500 switch.

When a group exists in the MLD snooping database, the switch responds to a group-specific query by sending an MLDv1 report. When the group is unknown, the group-specific query is flooded to the ingress VLAN.

When a host wants to leave a multicast group, it can send out an MLD Done message (equivalent to IGMP Leave message). When the switch receives an MLDv1 Done message, if Immediate-Leave is not enabled, the switch sends an MASQ to the port from which the message was received to determine if other devices connected to the port should remain in the multicast group.

Multicast Client Aging Robustness

You can configure port membership removal from addresses based on the number of queries. A port is removed from membership to an address only when there are no reports to the address on the port for the configured number of queries. The default number is 2.

Multicast Router Discovery

Like IGMP snooping, MLD snooping performs multicast router discovery, with these characteristics:

- Ports configured by a user never age out.
- Dynamic port learning results from MLDv1 snooping queries and IPv6 PIMv2 packets.
- If there are multiple routers on the same Layer 2 interface, MLD snooping tracks a single multicast router on the port (the router that most recently sent a router control packet).
- Dynamic multicast router port aging is based on a default timer of 5 minutes; the multicast router is deleted from the router port list if no control packet is received on the port for 5 minutes.
- IPv6 multicast router discovery only takes place when MLD snooping is enabled on the switch.
- Received IPv6 multicast router control packets are always flooded to the ingress VLAN, whether or not MLD snooping is enabled on the switch.
- After the discovery of the first IPv6 multicast router port, unknown IPv6 multicast data is forwarded only to the discovered router ports (before that time, all IPv6 multicast data is flooded to the ingress VLAN).

MLD Reports

The processing of MLDv1 join messages is essentially the same as with IGMPv2. When no IPv6 multicast routers are detected in a VLAN, reports are not processed or forwarded from the switch. When IPv6 multicast routers are detected and an MLDv1 report is received, an IPv6 multicast group address and an IPv6 multicast MAC address are entered in the VLAN MLD database. Then all IPv6 multicast traffic to the group within the VLAN is forwarded using this address. When MLD snooping is disabled, reports are flooded in the ingress VLAN.

When MLD snooping is enabled, MLD report suppression, called listener message suppression, is automatically enabled. With report suppression, the switch forwards the first MLDv1 report received by a group to IPv6 multicast routers; subsequent reports for the group are not sent to the routers. When MLD snooping is disabled, report suppression is disabled, and all MLDv1 reports are flooded to the ingress VLAN.

The switch also supports MLDv1 proxy reporting. When an MLDv1 MASQ is received, the switch responds with MLDv1 reports for the address on which the query arrived if the group exists in the switch on another port and if the port on which the query arrived is not the last member port for the address.

MLD Done Messages and Immediate-Leave

When the Immediate-Leave feature is enabled and a host sends an MLDv1 Done message (equivalent to an IGMP leave message), the port on which the Done message was received is immediately deleted from the group.You enable Immediate-Leave on VLANs and (as with IGMP snooping), you should only use the feature on VLANs where a single host is connected to the port. If the port was the last member of a group, the group is also deleted, and the leave information is forwarded to the detected IPv6 multicast routers.

When Immediate Leave is not enabled in a VLAN (which would be the case when there are multiple clients for a group on the same port) and a Done message is received on a port, an MASQ is generated on that port. The user can control when a port membership is removed for an existing address in terms of the number of MASQs. A port is removed from membership to an address when there are no MLDv1 reports to the address on the port for the configured number of queries.

The number of MASQs generated is configured by using the **ipv6 mld snooping last-listener-query count** global configuration command. The default number is 2.

The MASQ is sent to the IPv6 multicast address for which the Done message was sent. If there are no reports sent to the IPv6 multicast address specified in the MASQ during the switch maximum response time, the port on which the MASQ was sent is deleted from the IPv6 multicast address database. The maximum response time is the time configured by using the **ipv6 mld snooping last-listener-query-interval** global configuration command. If the deleted port is the last member of the multicast address, the multicast address is also deleted, and the switch sends the address leave information to all detected multicast routers.

Topology Change Notification Processing

When topology change notification (TCN) solicitation is enabled by using the **ipv6 mld snooping tcn query solicit** global configuration command, MLDv1 snooping sets the VLAN to flood all IPv6 multicast traffic with a configured number of MLDv1 queries before it begins sending multicast data only to selected ports. You set this value by using the **ipv6 mld snooping tcn flood query count** global configuration command. The default is to send two queries. The switch also generates MLDv1 global Done messages with valid link-local IPv6 source addresses when the switch becomes the STP root in the VLAN or when it is configured by the user. This is same as done in IGMP snooping.

Default MLD Snooping Configuration

Table 44-1 Default MLD Snooping Configuration

Feature	Default Setting
MLD snooping (Global)	Disabled.
MLD snooping (per VLAN)	Enabled. MLD snooping must be globally enabled for VLAN MLD snooping to take place.
IPv6 Multicast addresses	None configured.
IPv6 Multicast router ports	None configured.
MLD snooping Immediate Leave	Disabled.
MLD snooping robustness variable	Global: 2; Per VLAN: 0. **Note** The VLAN value overrides the global setting. When the VLAN value is 0, the VLAN uses the global count.
Last listener query count	Global: 2; Per VLAN: 0. **Note** The VLAN value overrides the global setting. When the VLAN value is 0, the VLAN uses the global count.
Last listener query interval	Global: 1000 (1 second); VLAN: 0. **Note** The VLAN value overrides the global setting. When the VLAN value is 0, the VLAN uses the global interval.
TCN query solicit	Disabled.
TCN query count	2.
MLD listener suppression	Enabled.

MLD Snooping Configuration Guidelines

When configuring MLD snooping, consider these guidelines:

- You can configure MLD snooping characteristics at any time, but you must globally enable MLD snooping by using the **ipv6 mld snooping** global configuration command for the configuration to take effect.
- When the IPv6 multicast router is a Catalyst 6500 switch and you are using extended VLANs (in the range 1006 to 4096), IPv6 MLD snooping must be enabled on the extended VLAN on the Catalyst 6500 switch in order for the switch to receive queries on the VLAN. For normal-range VLANs (1 to 1005), it is not necessary to enable IPv6 MLD snooping on the VLAN on the Catalyst 6500 switch.
- MLD snooping and IGMP snooping act independently of each other. You can enable both features at the same time on the switch.
- The maximum number of multicast entries allowed on the switch is determined by the configured SDM template.
- The maximum number of address entries allowed for the switch is 1000.

Enabling or Disabling MLD Snooping

By default, IPv6 MLD snooping is globally disabled on the switch and enabled on all VLANs. When MLD snooping is globally disabled, it is also disabled on all VLANs. When you globally enable MLD snooping, the VLAN configuration overrides the global configuration. That is, MLD snooping is enabled only on VLAN interfaces in the default state (enabled).

You can enable and disable MLD snooping on a per-VLAN basis or for a range of VLANs, but if you globally disable MLD snooping, it is disabled in all VLANs. If global snooping is enabled, you can enable or disable VLAN snooping.

Multicast Router Port

Although MLD snooping learns about router ports through MLD queries and PIMv6 queries, you can also use the command-line interface (CLI) to add a multicast router port to a VLAN. To add a multicast router port (add a static connection to a multicast router), use the **ipv6 mld snooping vlan mrouter** global configuration command on the switch.

MLD Immediate Leave

When you enable MLDv1 Immediate Leave, the switch immediately removes a port from a multicast group when it detects an MLD Done message on that port. You should only use the Immediate-Leave feature when there is a single receiver present on every port in the VLAN. When there are multiple clients for a multicast group on the same port, you should not enable Immediate-Leave in a VLAN.

MLD Snooping Queries

When Immediate Leave is not enabled and a port receives an MLD Done message, the switch generates MASQs on the port and sends them to the IPv6 multicast address for which the Done message was sent. You can optionally configure the number of MASQs that are sent and the length of time the switch waits for a response before deleting the port from the multicast group.

How to Configure IPv6 MLD Snooping

Note When the IPv6 multicast router is a Catalyst 6500 switch and you are using extended VLANs (in the range 1006 to 4096), IPv6 MLD snooping must be enabled on the extended VLAN on the Catalyst 6500 switch in order for the switch to receive queries on the VLAN. For normal-range VLANs (1 to 1005), it is not necessary to enable IPv6 MLD snooping on the VLAN on the Catalyst 6500 switch.

Enabling or Disabling MLD Snooping

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ipv6 mld snooping**	Globally enables MLD snooping on the switch.

	Command	Purpose
Step 3	**ipv6 mld snooping vlan** *vlan-id*	(Optional) Enables MLD snooping on the VLAN.The VLAN ID range is 1 to 1001 and 1006 to 4096. MLD snooping must be globally enabled for VLAN snooping to be enabled.
Step 4	**end**	Returns to privileged EXEC mode.
Step 5	**reload**	Reloads the operating system.

Configuring a Static Multicast Group

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode
Step 2	**ipv6 mld snooping vlan** *vlan-id* **static** *ipv6_multicast_address* **interface** *interface-id*	Statically configures a multicast group with a Layer 2 port as a member of a multicast group: • *vlan-id* is the multicast group VLAN ID. The VLAN ID range is 1 to 1001 and 1006 to 4096. • *ipv6_multicast_address* is the 128-bit group IPv6 address. The address must be in the form specified in RFC 2373. • *interface-id* is the member port. It can be a physical interface or a port channel (1 to 48).
Step 3	**end**	Returns to privileged EXEC mode.

Configuring a Multicast Router Port

Note Static connections to multicast routers are supported only on switch ports.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ipv6 mld snooping vlan** *vlan-id* **mrouter interface** *interface-id*	Specifies the multicast router VLAN ID, and specifies the interface to the multicast router. • The VLAN ID range is 1 to 1001 and 1006 to 4096. • The interface can be a physical interface or a port channel. The port-channel range is 1 to 48.
Step 3	**end**	Returns to privileged EXEC mode.

Enabling MLD Immediate Leave

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ipv6 mld snooping vlan** *vlan-id* **immediate-leave**	Enables MLD Immediate Leave on the VLAN interface.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring MLD Snooping Queries

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ipv6 mld snooping robustness-variable** *value*	(Optional) Sets the number of queries that are sent before switch will deletes a listener (port) that does not respond to a general query. The range is 1 to 3; the default is 2.
Step 3	**ipv6 mld snooping vlan** *vlan-id* **robustness-variable** *value*	(Optional) Sets the robustness variable on a VLAN basis, which determines the number of general queries that MLD snooping sends before aging out a multicast address when there is no MLD report response. The range is 1 to 3; the default is 0. When set to 0, the number used is the global robustness variable value.
Step 4	**ipv6 mld snooping last-listener-query-count** *count*	(Optional) Sets the number of MASQs that the switch sends before aging out an MLD client. The range is 1 to 7; the default is 2. The queries are sent 1 second apart.
Step 5	**ipv6 mld snooping vlan** *vlan-id* **last-listener-query-count** *count*	(Optional) Sets the last-listener query count on a VLAN basis. This value overrides the value configured globally. The range is 1 to 7; the default is 0. When set to 0, the global count value is used. Queries are sent 1 second apart.
Step 6	**ipv6 mld snooping last-listener-query-interval** *interval*	(Optional) Sets the maximum response time that the switch waits after sending out a MASQ before deleting a port from the multicast group. The range is 100 to 32,768 thousands of a second. The default is 1000 (1 second).
Step 7	**ipv6 mld snooping vlan** *vlan-id* **last-listener-query-interval** *interval*	(Optional) Sets the last-listener query interval on a VLAN basis. This value overrides the value configured globally. The range is 0 to 32,768 thousands of a second. The default is 0. When set to 0, the global last-listener query interval is used.
Step 8	**ipv6 mld snooping tcn query solicit**	(Optional) Enables topology change notification (TCN) solicitation, which means that VLANs flood all IPv6 multicast traffic for the configured number of queries before sending multicast data to only those ports requesting to receive it. The default is for TCN to be disabled.
Step 9	**ipv6 mld snooping tcn flood query count** *count*	(Optional) When TCN is enabled, specifies the number of TCN queries to be sent. The range is from 1 to 10; the default is 2.
Step 10	**end**	Returns to privileged EXEC mode.

Disabling MLD Listener Message Suppression

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**no ipv6 mld snooping listener-message-suppression**	Disables MLD message suppression.
Step 3	**end**	Returns to privileged EXEC mode.

Monitoring and Maintaining IPv6 MLD Snooping

You can display MLD snooping information for dynamically learned and statically configured router ports and VLAN interfaces. You can also display MAC address multicast entries for a VLAN configured for MLD snooping.

Command	Purpose
show ipv6 mld snooping [**vlan** *vlan-id*]	Displays the MLD snooping configuration information for all VLANs on the switch or for a specified VLAN. (Optional) Enter **vlan** *vlan-id* to display information for a single VLAN. The VLAN ID range is 1 to 1001 and 1006 to 4096.
show ipv6 mld snooping mrouter [**vlan** *vlan-id*]	Displays information on dynamically learned and manually configured multicast router interfaces. When you enable MLD snooping, the switch automatically learns the interface to which a multicast router is connected. These are dynamically learned interfaces. (Optional) Enter **vlan** *vlan-id* to display information for a single VLAN. The VLAN ID range is 1 to 1001 and 1006 to 4096.
show ipv6 mld snooping querier [**vlan** *vlan-id*]	Displays information about the IPv6 address and incoming port for the most-recently received MLD query messages in the VLAN. (Optional) Enter **vlan** *vlan-id* to display information for a single VLAN. The VLAN ID range is 1 to 1001 and 1006 to 4096.
show ipv6 mld snooping multicast-address [**vlan** *vlan-id*] [**count** | **dynamic** | **user**]	Displays all IPv6 multicast address information or specific IPv6 multicast address information for the switch or a VLAN. • Enter **count** to show the group count on the switch or in a VLAN. • Enter **dynamic** to display MLD snooping learned group information for the switch or for a VLAN. • Enter **user** to display MLD snooping user-configured group information for the switch or for a VLAN.
show ipv6 mld snooping multicast-address vlan *vlan-id* [*ipv6-multicast-address*]	Displays MLD snooping for the specified VLAN and IPv6 multicast address.
show ipv6 mld snooping multicast-address user or **show ipv6 mld snooping multicast-address vlan** *vlan-id* **user**	Verifies the static member port and the IPv6 address.

Command	Purpose
show ipv6 mld snooping mrouter [**vlan** *vlan-id*]	Verifies that IPv6 MLD snooping is enabled on the VLAN interface.
show ipv6 mld snooping	Verifies that IPv6 MLD snooping report suppression is disabled.

Configuration Examples for Configuring IPv6 MLD Snooping

Statically Configure an IPv6 Multicast Group: Example

This example shows how to statically configure an IPv6 multicast group:

```
Switch# configure terminal
Switch(config)# ipv6 mld snooping vlan 2 static FF12::3 interface gigabitethernet1/1
Switch(config)# end
```

Adding a Multicast Router Port to a VLAN: Example

This example shows how to add a multicast router port to VLAN 200:

```
Switch# configure terminal
Switch(config)# ipv6 mld snooping vlan 200 mrouter interface gigabitethernet1/2
Switch(config)# exit
```

Enabling MLD Immediate Leave on a VLAN: Example

This example shows how to enable MLD Immediate Leave on VLAN 130:

```
Switch# configure terminal
Switch(config)# ipv6 mld snooping vlan 130 immediate-leave
Switch(config)# exit
```

Setting MLD Snooping Global Robustness: Example

This example shows how to set the MLD snooping global robustness variable to 3:

```
Switch# configure terminal
Switch(config)# ipv6 mld snooping robustness-variable 3
Switch(config)# exit
```

Setting MLD Snooping Last-Listener Query Parameters: Examples

This example shows how to set the MLD snooping last-listener query count for a VLAN to 3:

```
Switch# configure terminal
Switch(config)# ipv6 mld snooping vlan 200 last-listener-query-count 3
Switch(config)# exit
```

This example shows how to set the MLD snooping last-listener query interval (maximum response time) to 2000 (2 seconds):

```
Switch# configure terminal
```

```
Switch(config)# ipv6 mld snooping last-listener-query-interval 2000
Switch(config)# exit
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
SDM templates	Chapter 11, "Configuring SDM Templates."

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 45

Configuring Cisco IOS IP SLAs Operations

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for Configuring Cisco IOS IP SLAs Operations

- Before configuring any IP SLAs application, we recommend that you verify the operation type supported on your software image by using the **show ip sla application** privileged EXEC command.

Restrictions for Configuring Cisco IOS IP SLAs Operations

- The IP SLAs responder can be a Cisco IOS Layer 2, responder-configurable switch, such as a Catalyst 2960 or IE 2000 switch running the LAN Base image, or a Catalyst 3560 or 3750 switch running the IP base image. The responder does not need to support full IP SLAs functionality.
- The switch does not support Voice over IP (VoIP) service levels using the gatekeeper registration delay operations measurements. Before configuring any IP SLAs application, you can use the **show ip sla application** privileged EXEC command to verify that the operation type is supported on your software image.

Information About Configuring Cisco IOS IP SLAs Operations

This chapter describes how to use Cisco IOS IP Service Level Agreements (SLAs) on the switch. Cisco IP SLAs is a part of Cisco IOS software that allows Cisco customers to analyze IP service levels for IP applications and services by using active traffic monitoring—the generation of traffic in a continuous, reliable, and predictable manner—for measuring network performance. With Cisco IOS IP SLAs, service provider customers can measure and provide service level agreements, and enterprise customers can verify service levels, verify outsourced service level agreements, and understand network performance. Cisco IOS IP SLAs can perform network assessments, verify quality of service (QoS), ease the deployment of new services, and assist with network troubleshooting.

Cisco IOS IP SLAs

Cisco IOS IP SLAs sends data across the network to measure performance between multiple network locations or across multiple network paths. It simulates network data and IP services and collects network performance information in real time. Cisco IOS IP SLAs generates and analyzes traffic either between Cisco IOS devices or from a Cisco IOS device to a remote IP device such as a network application server. Measurements provided by the various Cisco IOS IP SLAs operations can be used for troubleshooting, for problem analysis, and for designing network topologies.

Depending on the specific Cisco IOS IP SLAs operation, various network performance statistics are monitored within the Cisco device and stored in both command-line interface (CLI) and Simple Network Management Protocol (SNMP) MIBs. IP SLAs packets have configurable IP and application layer options such as source and destination IP address, User Datagram Protocol (UDP)/TCP port numbers, a type of service (ToS) byte (including Differentiated Services Code Point [DSCP] and IP Prefix bits), Virtual Private Network (VPN) routing/forwarding instance (VRF), and URL web address.

Because Cisco IP SLAs is Layer 2 transport independent, you can configure end-to-end operations over disparate networks to best reflect the metrics that an end user is likely to experience. IP SLAs collects a unique subset of these performance metrics:

- Delay (both round-trip and one-way)
- Jitter (directional)
- Packet loss (directional)
- Packet sequencing (packet ordering)
- Path (per hop)
- Connectivity (directional)
- Server or website download time

Because Cisco IOS IP SLAs is SNMP-accessible, it can also be used by performance-monitoring applications like CiscoWorks Internetwork Performance Monitor (IPM) and other third-party Cisco partner performance management products. Using IP SLAs can provide these benefits:

- Service-level agreement monitoring, measurement, and verification.
- Network performance monitoring
  - Measures the jitter, latency, or packet loss in the network.
  - Provides continuous, reliable, and predictable measurements.
- IP service network health assessment to verify that the existing QoS is sufficient for new IP services.
- Edge-to-edge network availability monitoring for proactive verification and connectivity testing of network resources (for example, shows the network availability of an NFS server used to store business critical data from a remote site).
- Troubleshooting of network operation by providing consistent, reliable measurement that immediately identifies problems and saves troubleshooting time.
- Multiprotocol Label Switching (MPLS) performance monitoring and network verification (if the switch supports MPLS)

Cisco IOS IP SLAs to Measure Network Performance

You can use IP SLAs to monitor the performance between any area in the network—core, distribution, and edge—without deploying a physical probe. It uses generated traffic to measure network performance between two networking devices. Figure 45-1 shows how IP SLAs begins when the source device sends a generated packet to the destination device. After the destination device receives the packet, depending on the type of IP SLAs operation, it responds with time-stamp information for the source to make the calculation on performance metrics. An IP SLAs operation performs a network measurement from the source device to a destination in the network using a specific protocol such as UDP.

Figure 45-1 Cisco IOS IP SLAs Operation

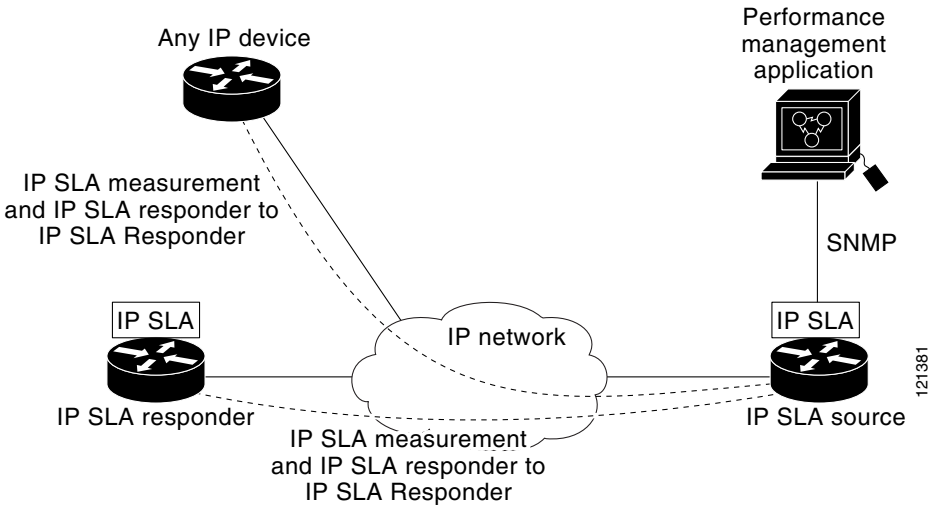

To implement IP SLAs network performance measurement, you need to perform these tasks:

1. Enable the IP SLAs responder, if required.
2. Configure the required IP SLAs operation type.
3. Configure any options available for the specified operation type.
4. Configure threshold conditions, if required.
5. Schedule the operation to run, then let the operation run for a period of time to gather statistics.
6. Display and interpret the results of the operation using the Cisco IOS CLI or a network management system (NMS) system with SNMP.

IP SLAs Responder and IP SLAs Control Protocol

The IP SLAs responder is a component embedded in the destination Cisco device that allows the system to anticipate and respond to IP SLAs request packets. The responder provides accurate measurements without the need for dedicated probes. The responder uses the Cisco IOS IP SLAs Control Protocol to provide a mechanism through which it can be notified on which port it should listen and respond. Only a Cisco IOS device can be a source for a destination IP SLAs Responder.

Figure 45-1 shows where the Cisco IOS IP SLAs responder fits in the IP network. The responder listens on a specific port for control protocol messages sent by an IP SLAs operation. Upon receipt of the control message, it enables the specified UDP or TCP port for the specified duration. During this time,

the responder accepts the requests and responds to them. It disables the port after it responds to the IP SLAs packet, or when the specified time expires. MD5 authentication for control messages is available for added security.

You do not need to enable the responder on the destination device for all IP SLAs operations. For example, a responder is not required for services that are already provided by the destination router (such as Telnet or HTTP). You cannot configure the IP SLAs responder on non-Cisco devices and Cisco IOS IP SLAs can send operational packets only to services native to those devices.

Response Time Computation for IP SLAs

Switches and routers can take tens of milliseconds to process incoming packets due to other high priority processes. This delay affects the response times because the test-packet reply might be in a queue while waiting to be processed. In this situation, the response times would not accurately represent true network delays. IP SLAs minimizes these processing delays on the source device as well as on the target device (if the responder is being used) to determine true round-trip times. IP SLAs test packets use time stamping to minimize the processing delays.

When the IP SLAs responder is enabled, it allows the target device to take time stamps when the packet arrives on the interface at interrupt level and again just as it is leaving, eliminating the processing time. This time stamping is made with a granularity of sub-milliseconds (ms).

Figure 45-2 demonstrates how the responder works. Four time stamps are taken to make the calculation for round-trip time. At the target router, with the responder functionality enabled, time stamp 2 (TS2) is subtracted from time stamp 3 (TS3) to produce the time spent processing the test packet as represented by delta. This delta value is then subtracted from the overall round-trip time. Notice that the same principle is applied by IP SLAs on the source router where the incoming time stamp 4 (TS4) is also taken at the interrupt level to allow for greater accuracy.

Figure 45-2 Cisco IOS IP SLAs Responder Time Stamping

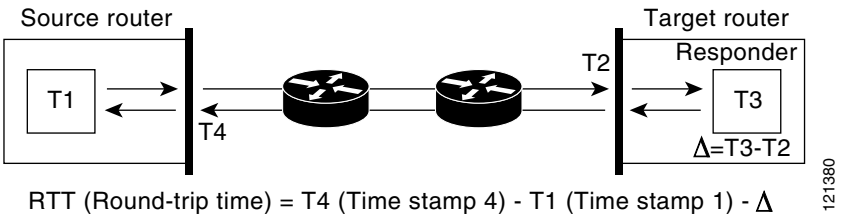

An additional benefit of the two time stamps at the target device is the ability to track one-way delay, jitter, and directional packet loss. Because much network behavior is asynchronous, it is critical to have these statistics. However, to capture one-way delay measurements, you must configure both the source router and target router with Network Time Protocol (NTP) so that the source and target are synchronized to the same clock source. One-way jitter measurements do not require clock synchronization.

IP SLAs Operation Scheduling

When you configure an IP SLAs operation, you must schedule the operation to begin capturing statistics and collecting error information. You can schedule an operation to start immediately or to start at a certain month, day, and hour. You can use the pending option to set the operation to start at a later time.

The pending option is an internal state of the operation that is visible through SNMP. The pending state is also used when an operation is a reaction (threshold) operation waiting to be triggered. You can schedule a single IP SLAs operation or a group of operations at one time.

You can schedule several IP SLAs operations on a switch running the IP services image by using a single command through the Cisco IOS CLI or the CISCO RTTMON-MIB. Scheduling the operations to run at evenly distributed times allows you to control the amount of IP SLAs monitoring traffic. This distribution of IP SLAs operations helps minimize the CPU utilization and thus improves network scalability.

IP SLAs Operation Threshold Monitoring

To support successful service level agreement monitoring, you must have mechanisms that notify you immediately of any possible violation. IP SLAs can send SNMP traps that are triggered by events such as these:

- Connection loss
- Timeout
- Round-trip time threshold
- Average jitter threshold
- One-way packet loss
- One-way jitter
- One-way mean opinion score (MOS)
- One-way latency

An IP SLAs threshold violation can also trigger another IP SLAs operation for further analysis. For example, the frequency could be increased or an ICMP path echo or ICMP path jitter operation could be initiated for troubleshooting.

Determining the type of threshold and the level to set can be complex, and depends on the type of IP service being used in the network.

IP Service Levels by Using the UDP Jitter Operation

Jitter means interpacket delay variance. When multiple packets are sent consecutively 10 ms apart from source to destination, if the network is behaving correctly, the destination should receive them 10 ms apart. But if there are delays in the network (like queuing, arriving through alternate routes, and so on) the arrival delay between packets might be more than or less than 10 ms with a positive jitter value meaning that the packets arrived more than 10 ms apart. If the packets arrive 12 ms apart, positive jitter is 2 ms; if the packets arrive 8 ms apart, negative jitter is 2 ms. For delay-sensitive networks, positive jitter values are undesirable, and a jitter value of 0 is ideal.

In addition to monitoring jitter, the IP SLAs UDP jitter operation can be used as a multipurpose data gathering operation. The packets IP SLAs generates carry packet sending and receiving sequence information and sending and receiving time stamps from the source and the operational target. Based on these, UDP jitter operations measure this data:

- Per-direction jitter (source to destination and destination to source)
- Per-direction packet-loss
- Per-direction delay (one-way delay)

- Round-trip delay (average round-trip time)

Because the paths for the sending and receiving of data can be different (asymmetric), you can use the per-direction data to more readily identify where congestion or other problems are occurring in the network.

The UDP jitter operation generates synthetic (simulated) UDP traffic and sends a number of UDP packets, each of a specified size, sent a specified number of milliseconds apart, from a source router to a target router, at a given frequency. By default, ten packet-frames, each with a payload size of 10 bytes are generated every 10 ms, and the operation is repeated every 60 seconds. You can configure each of these parameters to best simulate the IP service you want to provide.

To provide accurate one-way delay (latency) measurements, time synchronization, such as that provided by NTP, is required between the source and the target device. Time synchronization is not required for the one-way jitter and packet loss measurements. If the time is not synchronized between the source and target devices, one-way jitter and packet loss data is returned, but values of *0* are returned for the one-way delay measurements provided by the UDP jitter operation

Note Before you configure a UDP jitter operation on the source device, you must enable the IP SLAs responder on the target device (the operational target).

IP Service Levels by Using the ICMP Echo Operation

The ICMP echo operation measures end-to-end response time between a Cisco device and any devices using IP. Response time is computed by measuring the time taken between sending an ICMP echo request message to the destination and receiving an ICMP echo reply. Many customers use IP SLAs ICMP-based operations, in-house ping testing, or ping-based dedicated probes for response time measurements between the source IP SLAs device and the destination IP device. The IP SLAs ICMP echo operation conforms to the same specifications as ICMP ping testing, and the two methods result in the same response times.

Note This operation does not require the IP SLAs responder to be enabled.

How to Configure Cisco IOS IP SLAs Operations

Note Not all of the IP SLAs commands or operations described in this guide are supported on the switch. The switch supports IP service level analysis by using UDP jitter, UDP echo, HTTP, TCP connect, ICMP echo, ICMP path echo, ICMP path jitter, FTP, DNS, and DHCP, as well as multiple operation scheduling and proactive threshold monitoring. It does not support VoIP service levels using the gatekeeper registration delay operations measurements.

Configuring the IP SLAs Responder

Before You Begin

For the IP SLAs responder to function, you must also configure a source device, such as a Catalyst 3750 or Catalyst 3560 switch running the IP services image, that has full IP SLAs support. Refer to the documentation for the source device for configuration information.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip sla responder** {**tcp-connect** | **udp-echo**} **ipaddress** *ip-address* **port** *port-number*	Configures the switch as an IP SLAs responder. The optional keywords have these meanings: • **tcp-connect**—Enables the responder for TCP connect operations. • **udp-echo**—Enables the responder for User Datagram Protocol (UDP) echo or jitter operations. • **ipaddress** *ip-address*—Enters the destination IP address. • **port** *port-number*—Enters the destination port number. **Note** The IP address and port number must match those configured on the source device for the IP SLAs operation.
Step 3	**end**	Returns to privileged EXEC mode.

Configuring UDP Jitter Operation

Before You Begin

Before you configure a UDP jitter operation on the source device, you must enable the IP SLAs responder on the target device (the operational target).

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip sla** *operation-number*	Creates an IP SLAs operation, and enters IP SLAs configuration mode.

	Command	Purpose
Step 3	**udp-jitter** {*destination-ip-address* | *destination-hostname*} *destination-port* [**source-ip** {*ip-address* | *hostname*}] [**source-port** *port-number*] [**control** {**enable** | **disable**}] [**num-packets** *number-of-packets*] [**interval** *interpacket-interval*]	Configures the IP SLAs operation as a UDP jitter operation, and enters UDP jitter configuration mode. • *destination-ip-address* | *destination-hostname*—Specifies the destination IP address or hostname. • *destination-port*—Specifies the destination port number in the range from 1 to 65535. • (Optional) **source-ip** {*ip-address* | *hostname*}—Specifies the source IP address or hostname. When a source IP address or hostname is not specified, IP SLAs chooses the IP address nearest to the destination. • (Optional) **source-port** *port-number*—Specifies the source port number in the range from 1 to 65535. When a port number is not specified, IP SLAs chooses an available port. • (Optional) **control**—Enables or disables sending of IP SLAs control messages to the IP SLAs responder. By default, IP SLAs control messages are sent to the destination device to establish a connection with the IP SLAs responder. • (Optional) **num-packets** *number-of-packets*—Enters the number of packets to be generated. The range is 1 to 6000; the default is 10. • (Optional) **interval** *inter-packet-interval*—Enters the interval between sending packets in milliseconds. The range is 1 to 6000; the default value is 20 ms.
Step 4	**frequency** *seconds*	(Optional) Sets the rate at which a specified IP SLAs operation repeats. The range is from 1 to 604800 seconds; the default is 60 seconds.
Step 5	**exit**	Exits UDP jitter configuration mode, and returns to global configuration mode.

	Command	Purpose
Step 6	**ip sla schedule** *operation-number* [**life** {**forever** | *seconds*}] [**start-time** {*hh*:*mm* [:*ss*] [*month day* | *day month*] | **pending** | **now** | **after** *hh:mm:ss*] [**ageout** *seconds*] [**recurring**]	Configures the scheduling parameters for an individual IP SLAs operation. • *operation-number*—Enters the RTR entry number. • (Optional) **life**—Sets the operation to run indefinitely (**forever**) or for a specific number of *seconds*. The range is from 0 to 2147483647. The default is 3600 seconds (1 hour). • (Optional) **start-time**—Enters the time for the operation to begin collecting information: – To start at a specific time, enter the hour, minute, second (in 24-hour notation), and day of the month. If no month is entered, the default is the current month. – Enter **pending** to select no information collection until a start time is selected. – Enter **now** to start the operation immediately. – Enter **after** *hh:mm:ss* to show that the operation should start after the entered time has elapsed. • (Optional) **ageout** *seconds*—Enters the number of seconds to keep the operation in memory when it is not actively collecting information. The range is 0 to 2073600 seconds, the default is 0 seconds (never ages out). • (Optional) **recurring**—Sets the operation to automatically run every day.
Step 7	**end**	Returns to privileged EXEC mode.

Analyzing IP Service Levels by Using the ICMP Echo Operation

Note This operation does not require the IP SLAs responder to be enabled.

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**ip sla** *operation-number*	Creates an IP SLAs operation and enters IP SLAs configuration mode.
Step 3	**icmp-echo** {*destination-ip-address* | *destination-hostname*} [**source-ip** {*ip-address* | *hostname*} | **source-interface** *interface-id*]	Configures the IP SLAs operation as an ICMP Echo operation and enters ICMP echo configuration mode. • *destination-ip-address* | *destination-hostname*—Specifies the destination IP address or hostname. • (Optional) **source-ip** {*ip-address* | *hostname*}—Specifies the source IP address or hostname. When a source IP address or hostname is not specified, IP SLAs chooses the IP address nearest to the destination. • (Optional) **source-interface** *interface-id*—Specifies the source interface for the operation.
Step 4	**frequency** *seconds*	(Optional) Sets the rate at which a specified IP SLAs operation repeats. The range is from 1 to 604800 seconds; the default is 60 seconds.

	Command	Purpose
Step 5	**exit**	Exits UDP jitter configuration mode, and returns to global configuration mode.
Step 6	**ip sla schedule** *operation-number* [**life** {**forever** | *seconds*}] [**start-time** {*hh*:*mm* [:*ss*] [*month day* | *day month*] | **pending** | **now** | **after** *hh:mm:ss*] [**ageout** *seconds*] [**recurring**]	Configures the scheduling parameters for an individual IP SLAs operation. • *operation-number*—Enters the RTR entry number. • (Optional) **life**—Sets the operation to run indefinitely (**forever**) or for a specific number of *seconds*. The range is from 0 to 2147483647. The default is 3600 seconds (1 hour). • (Optional) **start-time**—Enters the time for the operation to begin collecting information: – To start at a specific time, enter the hour, minute, second (in 24-hour notation), and day of the month. If no month is entered, the default is the current month. – Enter **pending** to select no information collection until a start time is selected. – Enter **now** to start the operation immediately. – Enter **after** *hh:mm:ss* to indicate that the operation should start after the entered time has elapsed. • (Optional) **ageout** *seconds*—Enters the number of seconds to keep the operation in memory when it is not actively collecting information. The range is 0 to 2073600 seconds; the default is 0 seconds (never ages out). • (Optional) **recurring**—Sets the operation to automatically run every day.
Step 7	end	Returns to privileged EXEC mode.

Monitoring and Maintaining Cisco IP SLAs Operations

Command	Purpose
show ip sla application	Displays global information about Cisco IOS IP SLAs.
show ip sla authentication	Displays IP SLAs authentication information.
show ip sla configuration [*entry-number*]	Displays configuration values including all defaults for all IP SLAs operations or a specific operation.
show ip sla enhanced-history {**collection-statistics** | **distribution statistics**} [*entry-number*]	Displays enhanced history statistics for collected history buckets or distribution statistics for all IP SLAs operations or a specific operation.
show ip sla ethernet-monitor configuration [*entry-number*]	Displays IP SLAs automatic Ethernet configuration.
show ip sla event-publisher	Displays the list of client applications that are registered to receive IP SLAs notifications.
show ip sla group schedule [*schedule-entry-number*]	Displays IP SLAs group scheduling configuration and details.
show ip sla history [*entry-number* | **full** | **tabular**]	Displays history collected for all IP SLAs operations

Command	Purpose
show ip sla mpls-lsp-monitor {**collection-statistics** | **configuration** | **ldp operational-state** | **scan-queue** | **summary** [*entry-number*] | **neighbors**}	Displays MPLS label switched path (LSP) Health Monitor operations.
show ip sla reaction-configuration [*entry-number*]	Displays the configured proactive threshold monitoring settings for all IP SLAs operations or a specific operation.
show ip sla reaction-trigger [*entry-number*]	Displays the reaction trigger information for all IP SLAs operations or a specific operation.
show ip sla responder	Displays information about the IP SLAs responder.
show ip sla standards	Displays information about the IP SLAs standards.
show ip sla statistics [*entry-number* | **aggregated** | **details**]	Displays current or aggregated operational status and statistics.

Configuration Examples for Configuring Cisco IP SLAs Operations

Configuring an ICMP Echo IP SLAs Operation: Example

This example shows how to configure an ICMP echo IP SLAs operation:

```
Switch(config)# ip sla 12
Switch(config-ip-sla)# icmp-echo 172.29.139.134
Switch(config-ip-sla-echo)# frequency 30
Switch(config-ip-sla-echo)# exit
Switch(config)# ip sla schedule 5 start-time now life forever
Switch(config)# end
Switch# show ip sla configuration 22
IP SLAs, Infrastructure Engine-II.

Entry number: 12
Owner:
Tag:
Type of operation to perform: echo
Target address: 2.2.2.2
Source address: 0.0.0.0
Request size (ARR data portion): 28
Operation timeout (milliseconds): 5000
Type Of Service parameters: 0x0
Verify data: No
Vrf Name:
Schedule:
    Operation frequency (seconds): 60
    Next Scheduled Start Time: Pending trigger
    Group Scheduled : FALSE
    Randomly Scheduled : FALSE
    Life (seconds): 3600
    Entry Ageout (seconds): never
    Recurring (Starting Everyday): FALSE
    Status of entry (SNMP RowStatus): notInService
Threshold (milliseconds): 5000
Distribution Statistics:
```

```
    Number of statistic hours kept: 2
    Number of statistic distribution buckets kept: 1
    Statistic distribution interval (milliseconds): 20
History Statistics:
    Number of history Lives kept: 0
    Number of history Buckets kept: 15
    History Filter Type: None
Enhanced History:
```

Sample Output for Show IP SLA Command: Example

This is an example of the output from the command:

```
Switch# show ip sla application

        IP SLAs
Version: 2.2.0 Round Trip Time MIB, Infrastructure Engine-II
Time of last change in whole IP SLAs: 22:17:39.117 UTC Fri Jun
Estimated system max number of entries: 15801

Estimated number of configurable operations: 15801
Number of Entries configured  : 0
Number of active Entries      : 0
Number of pending Entries     : 0
Number of inactive Entries    : 0

        Supported Operation Types
Type of Operation to Perform: 802.1agEcho
Type of Operation to Perform: 802.1agJitter
Type of Operation to Perform: dhcp
Type of Operation to Perform: dns
Type of Operation to Perform: echo
Type of Operation to Perform: ftp
Type of Operation to Perform: http
Type of Operation to Perform: jitter
Type of Operation to Perform: pathEcho
Type of Operation to Perform: pathJitter
Type of Operation to Perform: tcpConnect
Type of Operation to Perform: udpEcho

IP SLAs low memory water mark: 21741224
```

Configuring a Responder UDP Jitter IP SLAs Operation: Example

This example shows how to configure the device as a responder for the UDP jitter IP SLAs operation in the next procedure:

```
Switch(config)# ip sla responder udp-echo 172.29.139.134 5000
```

Configuring a UDP Jitter IP SLAs Operation: Example

This example shows how to configure a UDP jitter IP SLAs operation:

```
Switch(config)# ip sla 10
Switch(config-ip-sla)# udp-jitter 172.29.139.134 5000
Switch(config-ip-sla-jitter)# frequency 30
Switch(config-ip-sla-jitter)# exit
```

```
Switch(config)# ip sla schedule 5 start-time now life forever
Switch(config)# end
Switch# show ip sla configuration 10
IP SLAs, Infrastructure Engine-II.

Entry number: 10
Owner:
Tag:
Type of operation to perform: udp-jitter
Target address/Source address: 1.1.1.1/0.0.0.0
Target port/Source port: 2/0
Request size (ARR data portion): 32
Operation timeout (milliseconds): 5000
Packet Interval (milliseconds)/Number of packets: 20/10
Type Of Service parameters: 0x0
Verify data: No
Vrf Name:
Control Packets: enabled
Schedule:
    Operation frequency (seconds): 30
    Next Scheduled Start Time: Pending trigger
    Group Scheduled : FALSE
    Randomly Scheduled : FALSE
    Life (seconds): 3600
    Entry Ageout (seconds): never
    Recurring (Starting Everyday): FALSE
    Status of entry (SNMP RowStatus): notInService
Threshold (milliseconds): 5000
Distribution Statistics:
    Number of statistic hours kept: 2
    Number of statistic distribution buckets kept: 1
    Statistic distribution interval (milliseconds): 20
Enhanced History:
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, Release 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
IP SLAs commands and configuration	*Cisco IOS IP SLAs Configuration Guide* on Cisco.com *Cisco IOS IP SLAs Command Reference* on Cisco.com

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

CHAPTER 46

Troubleshooting

This chapter describes how to identify and resolve software problems related to the Cisco IOS software on the switch. Depending on the nature of the problem, you can use the command-line interface (CLI), Device Manager, or Network Assistant to identify and solve problems.

For additional troubleshooting information, such as LED descriptions, see the *Cisco IE 2000 Switch Hardware Installation Guide*.

Finding Feature Information

Your software release may not support all the features documented in this chapter. For the latest feature information and caveats, see the release notes for your platform and software release.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Information for Troubleshooting

Autonegotiation Mismatches Prevention

The IEEE 802.3ab autonegotiation protocol manages the switch settings for speed (10 Mb/s, 100 Mb/s, and 1000 Mb/s, excluding SFP module ports) and duplex (half or full). There are situations when this protocol can incorrectly align these settings, reducing performance. A mismatch occurs under these circumstances:

- A manually set speed or duplex parameter is different from the manually set speed or duplex parameter on the connected port.
- A port is set to autonegotiate, and the connected port is set to full duplex with no autonegotiation.

To maximize switch performance and ensure a link, follow one of these guidelines when changing the settings for duplex and speed:

- Let both ports autonegotiate both speed and duplex.
- Manually set the speed and duplex parameters for the ports on both ends of the connection.

Note If a remote device does not autonegotiate, configure the duplex settings on the two ports to match. The speed parameter can adjust itself even if the connected port does not autonegotiate.

SFP Module Security and Identification

Cisco small form-factor pluggable (SFP) modules have a serial EEPROM that contains the module serial number, the vendor name and ID, a unique security code, and cyclic redundancy check (CRC). When an SFP module is inserted in the switch, the switch software reads the EEPROM to verify the serial number, vendor name and vendor ID, and recompute the security code and CRC. If the serial number, the vendor name or vendor ID, the security code, or CRC is invalid, the software generates a security error message and places the interface in an error-disabled state.

Note The security error message references the GBIC_SECURITY facility. The switch supports SFP modules and does not support GBIC modules. Although the error message text refers to GBIC interfaces and modules, the security messages actually refer to the SFP modules and module interfaces. For more information about error messages, see the system message guide for this release.

If you are using a non-Cisco SFP module, remove the SFP module from the switch, and replace it with a Cisco module. After inserting a Cisco SFP module, use the **errdisable recovery cause gbic-invalid** global configuration command to verify the port status, and enter a time interval for recovering from the error-disabled state. After the elapsed interval, the switch brings the interface out of the error-disabled state and retries the operation. For more information about the **errdisable recovery** command, see the command reference for this release.

If the module is identified as a Cisco SFP module, but the system is unable to read vendor-data information to verify its accuracy, an SFP module error message is generated. In this case, you should remove and reinsert the SFP module. If it continues to fail, the SFP module might be defective.

Ping

The switch supports IP ping, which you can use to test connectivity to remote hosts. Ping sends an echo request packet to an address and waits for a reply. Ping returns one of these responses:

- Normal response—The normal response (*hostname* is alive) occurs in 1 to 10 seconds, depending on network traffic.
- Destination does not respond—If the host does not respond, a *no-answer* message is returned.
- Unknown host—If the host does not exist, an *unknown host* message is returned.
- Destination unreachable—If the default gateway cannot reach the specified network, a *destination-unreachable* message is returned.
- Network or host unreachable—If there is no entry in the route table for the host or network, a *network or host unreachable* message is returned.

Layer 2 Traceroute

The Layer 2 traceroute feature allows the switch to identify the physical path that a packet takes from a source device to a destination device. Layer 2 traceroute supports only unicast source and destination MAC addresses. It finds the path by using the MAC address tables of the switches in the path. When the switch detects a device in the path that does not support Layer 2 traceroute, the switch continues to send Layer 2 trace queries and lets them time out.

The switch can only identify the path from the source device to the destination device. It cannot identify the path that a packet takes from source host to the source device or from the destination device to the destination host.

Layer 2 Traceroute Usage Guidelines

- Cisco Discovery Protocol (CDP) must be enabled on all the devices in the network. For Layer 2 traceroute to function properly, do not disable CDP.

  If any devices in the physical path are transparent to CDP, the switch cannot identify the path through these devices. For more information about enabling CDP, see Chapter 32, "Configuring CDP."
- A switch is reachable from another switch when you can test connectivity by using the **ping** privileged EXEC command. All switches in the physical path must be reachable from each other.
- The maximum number of hops identified in the path is ten.
- You can enter the **traceroute mac** or the **traceroute mac ip** privileged EXEC command on a switch that is not in the physical path from the source device to the destination device. All switches in the path must be reachable from this switch.
- The **traceroute mac** command output shows the Layer 2 path only when the specified source and destination MAC addresses belong to the same VLAN. If you specify source and destination MAC addresses that belong to different VLANs, the Layer 2 path is not identified, and an error message appears.
- If you specify a multicast source or destination MAC address, the path is not identified, and an error message appears.
- If the source or destination MAC address belongs to multiple VLANs, you must specify the VLAN to which both the source and destination MAC addresses belong. If the VLAN is not specified, the path is not identified, and an error message appears.
- The **traceroute mac ip** command output shows the Layer 2 path when the specified source and destination IP addresses belong to the same subnet. When you specify the IP addresses, the switch uses the Address Resolution Protocol (ARP) to associate the IP addresses with the corresponding MAC addresses and the VLAN IDs.
  - If an ARP entry exists for the specified IP address, the switch uses the associated MAC address and identifies the physical path.
  - If an ARP entry does not exist, the switch sends an ARP query and tries to resolve the IP address. If the IP address is not resolved, the path is not identified, and an error message appears.

- When multiple devices are attached to one port through hubs (for example, multiple CDP neighbors are detected on a port), the Layer 2 traceroute feature is not supported. When more than one CDP neighbor is detected on a port, the Layer 2 path is not identified, and an error message appears.

IP Traceroute

You can use IP traceroute to identify the path that packets take through the network on a hop-by-hop basis. The command output displays all network layer (Layer 3) devices, such as routers, that the traffic passes through on the way to the destination.

Your switches can participate as the source or destination of the **traceroute** privileged EXEC command and might or might not appear as a hop in the **traceroute** command output. If the switch is the destination of the traceroute, it is displayed as the final destination in the traceroute output. Intermediate switches do not show up in the traceroute output if they are only bridging the packet from one port to another within the same VLAN. However, if the intermediate switch is a multilayer switch that is routing a particular packet, this switch shows up as a hop in the traceroute output.

The **traceroute** privileged EXEC command uses the Time To Live (TTL) field in the IP header to cause routers and servers to generate specific return messages. Traceroute starts by sending a User Datagram Protocol (UDP) datagram to the destination host with the TTL field set to 1. If a router finds a TTL value of 1 or 0, it drops the datagram and sends an Internet Control Message Protocol (ICMP) time-to-live-exceeded message to the sender. Traceroute finds the address of the first hop by examining the source address field of the ICMP time-to-live-exceeded message.

To identify the next hop, traceroute sends a UDP packet with a TTL value of 2. The first router decrements the TTL field by 1 and sends the datagram to the next router. The second router sees a TTL value of 1, discards the datagram, and returns the time-to-live-exceeded message to the source. This process continues until the TTL is incremented to a value large enough for the datagram to reach the destination host (or until the maximum TTL is reached).

To learn when a datagram reaches its destination, traceroute sets the UDP destination port number in the datagram to a very large value that the destination host is unlikely to be using. When a host receives a datagram destined to itself containing a destination port number that is unused locally, it sends an ICMP *port-unreachable* error to the source. Because all errors except port-unreachable errors come from intermediate hops, the receipt of a port-unreachable error means that this message was sent by the destination port.

TDR

You can use the Time Domain Reflector (TDR) feature to diagnose and resolve cabling problems. When running TDR, a local device sends a signal through a cable and compares the reflected signal to the initial signal.

TDR is supported only on 10/100 and 10/100/1000 copper Ethernet ports. It is not supported on SFP module ports.

TDR can detect these cabling problems:

- Open, broken, or cut twisted-pair wires—The wires are not connected to the wires from the remote device.
- Shorted twisted-pair wires—The wires are touching each other or the wires from the remote device. For example, a shorted twisted pair can occur if one wire of the twisted pair is soldered to the other wire.

If one of the twisted-pair wires is open, TDR can find the length at which the wire is open.

Use TDR to diagnose and resolve cabling problems in these situations:

- Replacing a switch
- Setting up a wiring closet
- Troubleshooting a connection between two devices when a link cannot be established or when it is not operating properly

Crashinfo Files

The crashinfo files save information that helps Cisco technical support representatives to debug problems that caused the Cisco IOS image to fail (crash). The switch writes the crash information to the console at the time of the failure. The switch creates two types of crashinfo files:

- Basic crashinfo file—The switch automatically creates this file the next time you boot up the Cisco IOS image after the failure.
- Extended crashinfo file—The switch automatically creates this file when the system is failing.

Basic crashinfo Files

The information in the basic file includes the Cisco IOS image name and version that failed, a list of the processor registers, and other switch-specific information. You can provide this information to the Cisco technical support representative by using the **show tech-support** privileged EXEC command.

Basic crashinfo files are kept in this directory on the flash file system:

flash:/crashinfo/.

The filenames are crashinfo_*n* where *n* is a sequence number.

Each new crashinfo file that is created uses a sequence number that is larger than any previously existing sequence number, so the file with the largest sequence number describes the most recent failure. Version numbers are used instead of a timestamp because the switches do not include a real-time clock. You cannot change the name of the file that the system will use when it creates the file. However, after the file is created, you can use the **rename** privileged EXEC command to rename it, but the contents of the renamed file will not be displayed by the **show tech-support** privileged EXEC command. You can delete crashinfo files by using the **delete** privileged EXEC command.

You can display the most recent basic crashinfo file (that is, the file with the highest sequence number at the end of its filename) by entering the **show tech-support** privileged EXEC command. You also can access the file by using any command that can copy or display files, such as the **more** or the **copy** privileged EXEC command.

Extended crashinfo Files

The switch creates the extended crashinfo file when the system is failing. The information in the extended file includes additional information that can help determine the cause of the switch failure. You provide this information to the Cisco technical support representative by manually accessing the file and using the **more** or the **copy** privileged EXEC command.

Extended crashinfo files are kept in this directory on the flash file system:

flash:/crashinfo_ext/.

The filenames are crashinfo_ext_*n* where *n* is a sequence number.

You can configure the switch to not create the extended creashinfo file by using the **no exception crashinfo** global configuration command.

CPU Utilization

This section lists some possible symptoms that could be caused by the CPU being too busy and shows how to verify a CPU utilization problem. Table 46-1 lists the primary types of CPU utilization problems that you can identify. It gives possible causes and corrective action with links to the *Troubleshooting High CPU Utilization* document on Cisco.com.

Excessive CPU utilization might result in these symptoms, but the symptoms could also result from other causes.

- Spanning tree topology changes
- EtherChannel links brought down due to loss of communication
- Failure to respond to management requests (ICMP ping, SNMP timeouts, slow Telnet or SSH sessions)
- UDLD flapping
- IP SLAs failures because of SLAs responses beyond an acceptable threshold
- DHCP or IEEE 802.1x failures if the switch does not forward or respond to requests

Problem and Cause for High CPU Utilization

To determine if high CPU utilization is a problem, enter the **show processes cpu sorted** privileged EXEC command. Note the underlined information in the first line of the output example.

```
Switch# show processes cpu sorted
CPU utilization for five seconds: 8%/0%; one minute: 7%; five minutes: 8%
PID Runtime(ms) Invoked uSecs 5Sec 1Min 5Min TTY Process
140 8820183 4942081 1784 0.63% 0.37% 0.30% 0 HRPC qos request
100 3427318 16150534 212 0.47% 0.14% 0.11% 0 HRPC pm-counters
192 3093252 14081112 219 0.31% 0.14% 0.11% 0 Spanning Tree
143 8 37 216 0.15% 0.01% 0.00% 0 Exec
...
<output truncated>
```

This example shows normal CPU utilization. The output shows that utilization for the last 5 seconds is *8%/0%*, which has this meaning:

- The total CPU utilization is 8 percent, including both time running Cisco IOS processes and time spent handling interrupts.
- The time spent handling interrupts is zero percent.

Table 46-1 Troubleshooting CPU Utilization Problems

Type of Problem	Cause	Corrective Action
Interrupt percentage value is almost as high as total CPU utilization value.	The CPU is receiving too many packets from the network.	Determine the source of the network packet. Stop the flow, or change the switch configuration. See the section on "Analyzing Network Traffic."
Total CPU utilization is greater than 50% with minimal time spent on interrupts.	One or more Cisco IOS process is consuming too much CPU time. This is usually triggered by an event that activated the process.	Identify the unusual event, and troubleshoot the root cause. See the section on "Debugging Active Processes."

- For complete information about CPU utilization and how to troubleshoot utilization problems, see the *Troubleshooting High CPU Utilization* document on Cisco.com.

How to Troubleshoot

Recovering from Software Failures

Switch software can be corrupted during an upgrade, by downloading the wrong file to the switch, and by deleting the image file. In all of these cases, the switch does not pass the power-on self-test (POST), and there is no connectivity.

This procedure uses the Xmodem Protocol to recover from a corrupt or wrong image file. There are many software packages that support the Xmodem Protocol, and this procedure is largely dependent on the emulation software that you are using.

This recovery procedure requires that you have physical access to the switch.

Step 1 From your PC, download the software image tar file (*image_filename.tar*) from Cisco.com.

The Cisco IOS image is stored as a bin file in a directory in the tar file. For information about locating the software image files on Cisco.com, see the release notes.

Step 2 Extract the bin file from the tar file.

- If you are using Windows, use a zip program that can read a tar file. Use the zip program to navigate to and extract the bin file.
- If you are using UNIX, follow these steps:

1. Display the contents of the tar file by using the **tar -tvf** <*image_filename.tar*> UNIX command.

```
switch% tar -tvf image_filename.tar
```

2. Locate the bin file, and extract it by using the **tar -xvf** <*image_filename.tar*> <*image_filename.bin*> UNIX command.

```
switch% tar -xvf image_filename.tar image_filename.binx

x image_name.bin, 3970586 bytes, 7756 tape blocks
```

3. Verify that the bin file was extracted by using the **ls -l** <*image_filename.bin*> UNIX command.

```
switch% ls -l image_filename.bin-rwxr-xr-x   1 bschuett eng      6365325 May 19
13:03
<insert path for lan base image>

-rw-r--r--   1 boba      3970586 Apr 21 12:00 image_name.bin
```

Step 3 Connect your PC with terminal-emulation software supporting the Xmodem Protocol to the switch console port.

Step 4 Set the line speed on the emulation software to 9600 baud.

Step 5 Unplug the switch power cord.

Step 6 Press the **Express Setup** button and at the same time, reconnect the power cord to the switch.

You can release the **Express Setup** button a second or two after the LED above port 1 goes off. Several lines of information about the software appear along with instructions:

```
The system has been interrupted prior to initializing the flash file system. The following
commands will initialize the flash file system, and finish loading the operating system
software#

flash_init
load_helper
boot
```

Step 7 Initialize the flash file system:

```
switch: flash_init
```

Step 8 If you had set the console port speed to anything other than 9600, it has been reset to that particular speed. Change the emulation software line speed to match that of the switch console port.

Step 9 Load any helper files:

```
switch: load_helper
```

Step 10 Start the file transfer by using the Xmodem Protocol.

```
switch: copy xmodem: flash:image_filename.bin
```

Step 11 After the Xmodem request appears, use the appropriate command on the terminal-emulation software to start the transfer and to copy the software image into flash memory.

Step 12 Boot the newly downloaded Cisco IOS image.

```
switch:boot flash:image_filename.bin
```

Step 13 Use the **archive download-sw** privileged EXEC command to download the software image to the switch.

Step 14 Use the **reload** privileged EXEC command to restart the switch and to verify that the new software image is operating properly.

Step 15 Delete the flash:*image_filename.bin* file from the switch.

Recovering from a Lost or Forgotten Password

If you lose or forget your password, you can delete the switch password and set a new one.

Before you begin, make sure that:

- You have physical access to the switch.
- At least one switch port is enabled and is not connected to a device.

To delete the switch password and set a new one, follow these steps:

Step 1 Press the **Express Setup** button until the SETUP LED blinks green and the LED of an available switch downlink port blinks green.

If no switch downlink port is available for your PC or laptop connection, disconnect a device from one of the switch downlink ports. Press the **Express Setup** button again until the SETUP LED and the port LED blink green.

Step 2 Connect your PC or laptop to the port with the blinking green LED.

The SETUP LED and the switch downlink port LED stop blinking and stay solid green.

Step 3 Press and hold the **Express Setup** button. Notice that the SETUP LED starts blinking green again. Continue holding the button until the SETUP LED turns solid green (approximately 5 seconds). Release the **Express Setup** button immediately.

This procedure deletes the password without affecting any other configuration settings. You can now access the switch without a password through the console port or by using Device Manager.

Step 4 Enter a new password through the device manager by using the Express Setup window or through the command line interface by using the **enable secret** global configuration command.

Recovering from Lost Cluster Member Connectivity

Some configurations can prevent the command switch from maintaining contact with member switches. If you are unable to maintain management contact with a member, and the member switch is forwarding packets normally, check for these conflicts:

- A member switch (Catalyst 3750, Catalyst 3560, Catalyst 3550, Catalyst 3500 XL, Catalyst 2970, Catalyst 2960, Catalyst 2950, Catalyst 2900 XL, Catalyst 2820, and Catalyst 1900 switch) cannot connect to the command switch through a port that is defined as a network port.
- Catalyst 3500 XL, Catalyst 2900 XL, Catalyst 2820, and Catalyst 1900 member switches must connect to the command switch through a port that belongs to the same management VLAN.
- A member switch (Catalyst 3750, Catalyst 3560, Catalyst 3550, Catalyst 2970, Catalyst 2960, Catalyst 2950, Catalyst 3500 XL, Catalyst 2900 XL, Catalyst 2820, and Catalyst 1900 switch) connected to the command switch through a secured port can lose connectivity if the port is disabled because of a security violation.

Executing Ping

If you attempt to ping a host in a different IP subnetwork, you must define a static route to the network or have IP routing configured to route between those subnets. For more information, see Chapter 41, “Configuring Static IP Unicast Routing.”

IP routing is disabled by default on all switches. If you need to enable or configure IP routing, see Chapter 41, “Configuring Static IP Unicast Routing.”

Beginning in privileged EXEC mode, use this command to ping another device on the network from the switch:

Command	Purpose
ping ip *host* | *address*	Pings a remote host through IP or by supplying the hostname or network address.

Note Other protocol keywords are available with the **ping** command, but they are not supported in this release.

This example shows how to ping an IP host:

```
Switch# ping 172.20.52.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echoes to 172.20.52.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
Switch#
```

Table 46-2 describes the possible ping character output.

Table 46-2 Ping Output Display Characters

Character	Description
!	Each exclamation point means receipt of a reply.
.	Each period means the network server timed out while waiting for a reply.
U	A destination unreachable error PDU was received.
C	A congestion experienced packet was received.
I	User interrupted test.
?	Unknown packet type.
&	Packet lifetime exceeded.

To end a ping session, enter the escape sequence (**Ctrl-^ X** by default). Simultaneously press and release the **Ctrl**, **Shift**, and **6** keys and then press the **X** key.

Executing IP Traceroute

Beginning in privileged EXEC mode, follow this step to trace that the path packets take through the network:

Command	Purpose
traceroute ip *host*	Traces the path that packets take through the network.

Note Other protocol keywords are available with the **traceroute** privileged EXEC command, but they are not supported in this release.

This example shows how to perform a **traceroute** to an IP host:

```
Switch# traceroute ip 171.9.15.10

Type escape sequence to abort.
Tracing the route to 171.69.115.10

  1 172.2.52.1 0 msec 0 msec 4 msec
  2 172.2.1.203 12 msec 8 msec 0 msec
  3 171.9.16.6 4 msec 0 msec 0 msec
  4 171.9.4.5 0 msec 4 msec 0 msec
  5 171.9.121.34 0 msec 4 msec 4 msec
  6 171.9.15.9 120 msec 132 msec 128 msec
  7 171.9.15.10 132 msec 128 msec 128 msec
Switch#
```

The display shows the hop count, the IP address of the router, and the round-trip time in milliseconds for each of the three probes that are sent.

Table 46-3 Traceroute Output Display Characters

Character	Description
*	The probe timed out.
?	Unknown packet type.
A	Administratively unreachable. Usually, this output means that an access list is blocking traffic.
H	Host unreachable.
N	Network unreachable.
P	Protocol unreachable.
Q	Source quench.
U	Port unreachable.

To end a trace in progress, enter the escape sequence (**Ctrl-^ X** by default). Simultaneously press and release the **Ctrl**, **Shift**, and **6** keys and then press the **X** key.

Running TDR and Displaying the Results

To run TDR, enter the **test cable-diagnostics tdr interface** *interface-id* privileged EXEC command:

To display the results, enter the **show cable-diagnostics tdr interface** *interface-id* privileged EXEC command. For a description of the fields in the display, see the command reference for this release.

Enabling Debugging on a Specific Feature

Caution Because debugging output is assigned high priority in the CPU process, it can render the system unusable. For this reason, use **debug** commands only to troubleshoot specific problems or during troubleshooting sessions with Cisco technical support staff. It is best to use **debug** commands during periods of lower network traffic and fewer users. Debugging during these periods decreases the likelihood that increased **debug** command processing overhead will affect system use.

All **debug** commands are entered in privileged EXEC mode, and most **debug** commands take no arguments. For example, beginning in privileged EXEC mode, enter this command to enable the debugging for Switched Port Analyzer (SPAN):

```
Switch# debug span-session
```

The switch continues to generate output until you enter the **no** form of the command.

If you enable a **debug** command and no output appears, consider these possibilities:

- The switch might not be properly configured to generate the type of traffic you want to monitor. Use the **show running-config** command to check its configuration.
- Even if the switch is properly configured, it might not generate the type of traffic you want to monitor during the particular period that debugging is enabled. Depending on the feature you are debugging, you can use commands such as the TCP/IP **ping** command to generate network traffic.

To disable debugging of SPAN, enter this command in privileged EXEC mode:

```
Switch# no debug span-session
```

Alternately, in privileged EXEC mode, you can enter the **undebug** form of the command:

```
Switch# undebug span-session
```

To display the state of each debugging option, enter this command in privileged EXEC mode:

```
Switch# show debugging
```

Enabling All-System Diagnostics

Beginning in privileged EXEC mode, enter this command to enable all-system diagnostics:

```
Switch# debug all
```

Caution Because debugging output takes priority over other network traffic, and because the **debug all** privileged EXEC command generates more output than any other **debug** command, it can severely diminish switch performance or even render it unusable. In virtually all cases, it is best to use more specific **debug** commands.

The **no debug all** privileged EXEC command disables all diagnostic output. Using the **no debug all** command is a convenient way to ensure that you have not accidentally left any **debug** commands enabled.

Redirecting Debug and Error Message Output

By default, the network server sends the output from **debug** commands and system error messages to the console. If you use this default, you can use a virtual terminal connection to monitor debug output instead of connecting to the console port.

Possible destinations include the console, virtual terminals, internal buffer, and UNIX hosts running a syslog server. The syslog format is compatible with 4.3 Berkeley Standard Distribution (BSD) UNIX and its derivatives.

Note Be aware that the debugging destination you use affects system overhead. Logging messages to the console produces very high overhead, whereas logging messages to a virtual terminal produces less overhead. Logging messages to a syslog server produces even less, and logging to an internal buffer produces the least overhead of any method.

For more information about system message logging, see Chapter 35, "Configuring System Message Logging."

Monitoring Information

Physical Path

You can display the physical path that a packet takes from a source device to a destination device by using one of these privileged EXEC commands:

- **tracetroute mac** [**interface** *interface-id*] {*source-mac-address*} [**interface** *interface-id*] {*destination-mac-address*} [**vlan** *vlan-id*] [**detail**]
- **tracetroute mac ip** {*source-ip-address* | *source-hostname*}{*destination-ip-address* | *destination-hostname*} [**detail**]

For more information, see the command reference for this release.

SFP Module Status

You can check the physical or operational status of an SFP module by using the **show interfaces transceiver** privileged EXEC command. This command shows the operational status, such as the temperature and the current for an SFP module on a specific interface and the alarm status. You can also use the command to check the speed and the duplex settings on an SFP module. For more information, see the **show interfaces transceiver** command in the command reference for this release.

Troubleshooting Examples

show platform forward Command

The output from the **show platform forward** privileged EXEC command provides some useful information about the forwarding results if a packet entering an interface is sent through the system. Depending upon the parameters entered about the packet, the output provides lookup table results and port maps used to calculate forwarding destinations, bitmaps, and egress information.

Most of the information in the output from the command is useful mainly for technical support personnel, who have access to detailed information about the switch application-specific integrated circuits (ASICs). However, packet forwarding information can also be helpful in troubleshooting.

This is an example of the output from the **show platform forward** command on port 1 in VLAN 5 when the packet entering that port is addressed to unknown MAC addresses. The packet should be flooded to all other ports in VLAN 5.

```
Switch# show platform forward gigabitethernet1/1 vlan 5 1.1.1 2.2.2 ip 13.1.1.1 13.2.2.2
udp 10 20
Global Port Number:24, Asic Number:5
Src Real Vlan Id:5, Mapped Vlan Id:5

Ingress:
 Lookup                   Key-Used                    Index-Hit  A-Data
InptACL  40_0D020202_0D010101-00_40000014_000A0000     01FFA   03000000
L2Local  80_00050002_00020002-00_00000000_00000000     00C71   0000002B
Station Descriptor:02340000, DestIndex:0239, RewriteIndex:F005

=========================================
Egress:Asic 2, switch 1
Output Packets:

-----------------------------------------
Packet 1
 Lookup                   Key-Used                    Index-Hit  A-Data
OutptACL 50_0D020202_0D010101-00_40000014_000A0000     01FFE   03000000

Port       Vlan      SrcMac          DstMac    Cos  Dscpv
Gi1/1    0005 0001.0001.0001  0002.0002.0002

-----------------------------------------
Packet 2
 Lookup                   Key-Used                    Index-Hit  A-Data
OutptACL 50_0D020202_0D010101-00_40000014_000A0000     01FFE   03000000

Port       Vlan      SrcMac          DstMac    Cos  Dscpv
Gi1/1    0005 0001.0001.0001  0002.0002.0002

-----------------------------------------
<output truncated>
-----------------------------------------
Packet 10
 Lookup                   Key-Used                    Index-Hit  A-Data
OutptACL 50_0D020202_0D010101-00_40000014_000A0000     01FFE   03000000
Packet dropped due to failed DEJA_VU Check on Gi1/0/2
Packet dropped due to failed DEJA_VU Check on Gi1/2
```

This is an example of the output when the packet coming in on port 1 in VLAN 5 is sent to an address already learned on the VLAN on another port. It should be forwarded from the port on which the address was learned.

```
Switch# show platform forward gigabitethernet1/1 vlan 5 1.1.1 0009.43a8.0145 ip 13.1.1.1
13.2.2.2 udp 10 20
Global Port Number:24, Asic Number:5
Src Real Vlan Id:5, Mapped Vlan Id:5

Ingress:
 Lookup                    Key-Used                      Index-Hit  A-Data
InptACL  40_0D020202_0D010101-00_40000014_000A0000     01FFA   03000000
L2Local  80_00050009_43A80145-00_00000000_00000000     00086   02010197
Station Descriptor:F0050003, DestIndex:F005, RewriteIndex:0003

========================================
Egress:Asic 3, switch 1
Output Packets:

----------------------------------------
Packet 1
 Lookup                    Key-Used                      Index-Hit  A-Data
OutptACL 50_0D020202_0D010101-00_40000014_000A0000     01FFE   03000000

Port            Vlan      SrcMac          DstMac    Cos  Dscpv
interface-id    0005 0001.0001.0001  0009.43A8.0145
```

This is an example of the output when the packet coming in on port 1 in VLAN 5 has a destination MAC address set to the router MAC address in VLAN 5 and the destination IP address unknown. Because there is no default route set, the packet should be dropped.

```
Switch# show platform forward gigabitethernet1/1 vlan 5 1.1.1 03.e319.ee44 ip 13.1.1.1
13.2.2.2 udp 10 20
Global Port Number:24, Asic Number:5
Src Real Vlan Id:5, Mapped Vlan Id:5

Ingress:
 Lookup                    Key-Used                      Index-Hit  A-Data
InptACL  40_0D020202_0D010101-00_41000014_000A0000     01FFA   03000000
L3Local  00_00000000_00000000-90_00001400_0D020202     010F0   01880290
L3Scndr  12_0D020202_0D010101-00_40000014_000A0000     034E0   000C001D_00000000
Lookup Used:Secondary
Station Descriptor:02260000, DestIndex:0226, RewriteIndex:0000
```

This is an example of the output when the packet coming in on port 1 in VLAN 5 has a destination MAC address set to the router MAC address in VLAN 5 and the destination IP address set to an IP address that is in the IP routing table. It should be forwarded as specified in the routing table.

```
Switch# show platform forward gigabitethernet1/1 vlan 5 1.1.1 03.e319.ee44 ip 110.1.5.5
16.1.10.5
Global Port Number:24, Asic Number:5
Src Real Vlan Id:5, Mapped Vlan Id:5

Ingress:
 Lookup                    Key-Used                      Index-Hit  A-Data
InptACL  40_10010A05_0A010505-00_41000014_000A0000     01FFA   03000000
L3Local  00_00000000_00000000-90_00001400_10010A05     010F0   01880290
L3Scndr  12_10010A05_0A010505-00_40000014_000A0000     01D28   30090001_00000000
Lookup Used:Secondary
Station Descriptor:F0070007, DestIndex:F007, RewriteIndex:0007

========================================
Egress:Asic 3, switch 1
```

```
Output Packets:

----------------------------------------
Packet 1
 Lookup                   Key-Used                    Index-Hit  A-Data
OutptACL 50_10010A05_0A010505-00_40000014_000A0000     01FFE   03000000

Port       Vlan      SrcMac          DstMac    Cos  Dscpv
Gi1/2    0007 XXXX.XXXX.0246  0009.43A8.0147
```

Additional References

The following sections provide references related to switch administration:

Related Documents

Related Topic	Document Title
Cisco IE 2000 commands	*Cisco IE 2000 Switch Command Reference*, 15.0(1)EY
Cisco IOS basic commands	*Cisco IOS Configuration Fundamentals Command Reference*
Additional troubleshooting information	*Cisco IE 2000 Switch Hardware Installation Guide*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

APPENDIX A

Working with the Cisco IOS File System, Configuration Files, and Software Images

This appendix describes how to manipulate the switch flash file system, how to copy configuration files, and how to archive (upload and download) software images to a switch.

Note For complete syntax and usage information for the commands used in this chapter, see the switch command reference for this release and the *Cisco IOS Configuration Fundamentals Command Reference, Release 15.0* from the Cisco.com page.

Working with the Flash File System

The flash file system is a single flash device on which you can store files. It also provides several commands to help you manage software image and configuration files. The default flash file system on the switch is named *flash:*.

The switch has a removable compact flash card that stores the Cisco IOS software image and configuration files. Removing the compact flash card does not interrupt switch operation unless you need to reload the Cisco IOS software. However, if you remove the compact flash card, you do not have access to the flash file system, and any attempt to access it generates an error message.

Use the **show flash:** privileged EXEC command to display the compact flash file settings. For more information about the command, go to this URL:

http://www.cisco.com/en/US/docs/ios/12_2/configfun/command/reference/frf009.html#wp1018357

For information about how to remove or replace the compact flash memory card on the switch, see the *Cisco IE 2000 Hardware Installation Guide*.

Displaying Available File Systems

To display the available file systems on your switch, use the **show file systems** privileged EXEC command as shown in this example.

```
Switch# show file systems

File Systems:

     Size(b)     Free(b)      Type  Flags  Prefixes
           -           -    opaque     ro   bs:
```

```
*  134086656   117346304     flash     rw   flash:
           -           -    opaque     rw   system:
           -           -    opaque     rw   tmpsys:
      524288      518334     nvram     rw   nvram:
           -           -    opaque     ro   xmodem:
           -           -    opaque     ro   ymodem:
           -           -    opaque     rw   null:
           -           -    opaque     ro   tar:
           -           -   network     rw   tftp:
           -           -   network     rw   rcp:
           -           -   network     rw   http:
           -           -   network     rw   ftp:
           -           -   network     rw   scp:
           -           -   network     rw   https:
           -           -    opaque     ro   cns:

Switch#
```

Detecting an Unsupported SD Flash Memory Card

When the switch starts and detects an unsupported Secure Digital (SD) flash memory card, or when you insert an unsupported SD flash memory card while the switch is running, the following warning message is displayed:

```
WARNING: Non-IT SD flash detected. Use of this card during normal
         operation can impact and severely degrade performance of the system.
          Please use supported SD flash cards only.
```

To display information about the SD flash memory card on the screen, use the **show platform sdflash** privileged EXEC command.

This example shows an unsupported SD flash memory card:

```
Switch# show platform sdflash
SD Flash Manufacturer      : SMART MODULAR (ID=27h) - Non IT
         Size              : 485MB
         Serial number     : B01000A5
         Revision          : 2.0
         Manufacturing date: 12/2009
```

This example shows a supported SD flash memory card:

```
Switch# show platform sdflash
SD Flash Manufacturer      : SMART MODULAR (ID=27h)
         Size                          : 972MB
         Serial number         : 07000019
         Revision                   : 2.0
         Manufacturing date: 3/2010
```

Note When you enter the **show platform sdflash** privileged EXEC command, the name, date, and other fields that are displayed depend on the manufacturer of the SD flash memory card. However, if the SD flash memory card is unsupported, "Non IT" is displayed after the manufacturer's name.

Note The output of the **show platform sdflash** privileged EXEC command is also included in the **show tech-support** privileged EXEC command output.

SD Flash Memory Card LED

Table A-1 SD Flash Memory Card LED

Color	System Status
Off / blinking green	SD flash memory card transfer in progress.
Slow blinking amber	SD flash memory card is unsupported.
Fast blinking amber	SD flash memory card is not present.
Amber	Error accessing the SD flash memory card. Cisco IOS boot image cannot be found.
Green	SD flash memory card is functioning.

Setting the Default File System

Table A-2 show file systems Field Descriptions

Field	Value
Size(b)	Amount of memory in the file system in bytes.
Free(b)	Amount of free memory in the file system in bytes.
Type	Type of file system. **flash**—The file system is for a flash memory device. **nvram**—The file system is for a NVRAM device. **opaque**—The file system is a locally generated *pseudo* file system (for example, the *system*) or a download interface, such as brimux. **unknown**—The file system is an unknown type.

Table A-2 show file systems Field Descriptions (continued)

Field	Value
Flags	Permission for file system. **ro**—read-only. **rw**—read/write.\ **wo**—write-only.
Prefixes	Alias for file system. **flash:**—Flash file system. **nvram:**—NVRAM. **null:**—Null destination for copies. You can copy a remote file to null to find its size. **rcp:**—Remote Copy Protocol (RCP) network server. **system:**—Contains the system memory, including the running configuration. **tftp:**—TFTP network server. **xmodem:**—Obtain the file from a network machine by using the Xmodem protocol. **ymodem:**—Obtain the file from a network machine by using the Ymodem protocol.

You can specify the file system or directory that the system uses as the default file system by using the **cd** *filesystem:* privileged EXEC command. You can set the default file system to omit the *filesystem:* argument from related commands. For example, for all privileged EXEC commands that have the optional *filesystem:* argument, the system uses the file system specified by the **cd** command.

By default, the default file system is *flash:*.

You can display the current default file system as specified by the **cd** command by using the **pwd** privileged EXEC command.

Displaying Information About Files on a File System

You can view a list of the contents of a file system before manipulating its contents. For example, before copying a new configuration file to flash memory, you might want to verify that the file system does not already contain a configuration file with the same name. Similarly, before copying a flash configuration file to another location, you might want to verify its filename for use in another command.

To display information about files on a file system, use one of the privileged EXEC commands in Table A-3.

Table A-3 Commands for Displaying Information About Files

Command	Description
dir [**/all**] [*filesystem***:**][*filename*]	Display a list of files on a file system.
show file systems	Display more information about each of the files on a file system.
show file information *file-url*	Display information about a specific file.
show file descriptors	Display a list of open file descriptors. File descriptors are the internal representations of open files. You can use this command to see if another user has a file open.

Changing Directories and Displaying the Working Directory

Beginning in privileged EXEC mode, follow these steps to change directories and display the working directory:

	Command	Purpose
Step 1	**dir** *filesystem***:**	Displays the directories on the specified file system. For *filesystem***:**, use **flash:** for the system board flash device.
Step 2	**cd new_configs**	Changes to the directory of interest. The command example shows how to change to the directory named *new_configs*.
Step 3	**pwd**	Displays the working directory.

Creating and Removing Directories

Beginning in privileged EXEC mode, follow these steps to create and remove a directory:

	Command	Purpose
Step 1	**dir** *filesystem***:**	Displays the directories on the specified file system. For *filesystem***:**, use **flash:** for the system board flash device.
Step 2	**mkdir old_configs**	Creates a new directory. The command example shows how to create the directory named *old_configs*. Directory names are case sensitive. Directory names are limited to 45 characters between the slashes (/); the name cannot contain control characters, spaces, deletes, slashes, quotes, semicolons, or colons.
Step 3	**dir** *filesystem***:**	Verifies your entry.

To delete a directory with all its files and subdirectories, use the **delete /force /recursive** *filesystem***:/***file-url* privileged EXEC command.

Use the **/recursive** keyword to delete the named directory and all subdirectories and the files contained in it. Use the **/force** keyword to suppress the prompting that confirms a deletion of each file in the directory. You are prompted only once at the beginning of this deletion process. Use the **/force** and **/recursive** keywords for deleting old software images that were installed by using the **archive download-sw** command but are no longer needed.

For *filesystem*, use **flash:** for the system board flash device. For *file-url*, enter the name of the directory to be deleted. All the files in the directory and the directory are removed.

Caution When files and directories are deleted, their contents cannot be recovered.

Copying Files

To copy a file from a source to a destination, use the **copy** *source-url destination-url* privileged EXEC command. For the source and destination URLs, you can use **running-config** and **startup-config** keyword shortcuts. For example, the **copy running-config startup-config** command saves the currently running configuration file to the NVRAM section of flash memory to be used as the configuration during system initialization.

You can also copy from special file systems (**xmodem:**, **ymodem:**) as the source for the file from a network machine that uses the Xmodem or Ymodem protocol.

Network file system URLs include **ftp:**, **rcp:**, and **tftp:** and have these syntaxes:

- FTP—**ftp:**[[**//***username* [**:***password*]**@***location*]**/***directory*]**/***filename*
- RCP—**rcp:**[[**//***username***@***location*]**/***directory*]**/***filename*
- TFTP—**tftp:**[[**//***location*]**/***directory*]**/***filename*

Local writable file systems include flash:.

Some invalid combinations of source and destination exist. Specifically, you cannot copy these combinations:

- From a running configuration to a running configuration
- From a startup configuration to a startup configuration
- From a device to the same device (for example, the **copy flash: flash:** command is invalid)

For specific examples of using the **copy** command with configuration files, see the "Working with Configuration Files" section on page A-9.

To copy software images either by downloading a new version or by uploading the existing one, use the **archive download-sw** or the **archive upload-sw** privileged EXEC command. For more information, see the "Working with Software Images" section on page A-22.

Deleting Files

When you no longer need a file on a flash memory device, you can permanently delete it. To delete a file or directory from a specified flash device, use the **delete** [**/force**] [**/recursive**] [*filesystem***:**]*/file-url* privileged EXEC command.

Use the **/recursive** keyword for deleting a directory and all subdirectories and the files contained in it. Use the **/force** keyword to suppress the prompting that confirms a deletion of each file in the directory. You are prompted only once at the beginning of this deletion process. Use the **/force** and **/recursive** keywords for deleting old software images that were installed by using the **archive download-sw** command but are no longer needed.

If you omit the *filesystem***:** option, the switch uses the default device specified by the **cd** command. For *file-url*, you specify the path (directory) and the name of the file to be deleted.

When you attempt to delete any files, the system prompts you to confirm the deletion.

Caution When files are deleted, their contents cannot be recovered.

This example shows how to delete the file *myconfig* from the default flash memory device:

```
Switch# delete myconfig
```

Creating, Displaying, and Extracting tar Files

You can create a tar file and write files into it, list the files in a tar file, and extract the files from a tar file as described in the next sections.

Note Instead of using the **copy** privileged EXEC command or the **archive tar** privileged EXEC command, we recommend using the **archive download-sw** and **archive upload-sw** privileged EXEC commands to download and upload software image files.

Creating a tar File

To create a tar file and write files into it, use this privileged EXEC command:

archive tar /create *destination-url* **flash:/***file-url*

For *destination-url*, specify the destination URL alias for the local or network file system and the name of the tar file to create. These options are supported:

- For the local flash file system, the syntax is **flash:**
- For the FTP, the syntax is **ftp:**[[**//***username*[**:***password*]**@***location*]**/***directory***/***tar-filename***.tar**
- For the RCP, the syntax is **rcp:**[[**//***username***@***location*]**/***directory*]**/***tar-filename***.tar**
- For the TFTP, the syntax is **tftp:**[[**//***location*]**/***directory*]**/***tar-filename***.tar**

The *tar-filename***.tar** is the tar file to be created.

For **flash:/***file-url*, specify the location on the local flash file system from which the new tar file is created. You can also specify an optional list of files or directories within the source directory to write to the new tar file. If none are specified, all files and directories at this level are written to the newly created tar file.

This example shows how to create a tar file. This command writes the contents of the *new-configs* directory on the local flash device to a file named *saved.tar* on the TFTP server at 172.20.10.30:

```
Switch# archive tar /create tftp:172.20.10.30/saved.tar flash:/new-configs
```

Displaying the Contents of a tar File

To display the contents of a tar file on the screen, use this privileged EXEC command:

archive tar /table *source-url*

For *source-url*, specify the source URL alias for the local or network file system. These options are supported:

- For the local flash file system, the syntax is **flash:**
- For the FTP, the syntax is **ftp:**[[**//***username*[**:***password*]**@***location*]**/***directory*]**/***tar-filename***.tar**
- For the RCP, the syntax is **rcp:**[[**//***username***@***location*]**/***directory*]**/***tar-filename***.tar**

- For the TFTP, the syntax is
  tftp:[[*//location*]*/directory*]*/tar-filename*.**tar**

The *tar-filename*.**tar** is the tar file to display.

You can also limit the display of the files by specifying an optional list of files or directories after the tar file; then only those files appear. If none are specified, all files and directories appear.

This example shows how to display the contents of a switch tar file that is in flash memory:

```
Switch# archive tar /table flash:image-name.tar
image-name/ (directory)
image-name/html/ (directory)
image-name/html/file.html (0 bytes)
image-name/image-name.bin (610856 bytes)
image-name/info (219 bytes)
```

This example shows how to display only the */html* directory and its contents:

```
Switch# archive tar /table flash: image-name/html
cimage-name/html
cimage-name/html/ (directory)
cimage-name/html/const.htm (556 bytes)
cimage-name/html/xhome.htm (9373 bytes)
cimage-name/html/menu.css (1654 bytes)
<output truncated>
```

Extracting a tar File

To extract a tar file into a directory on the flash file system, use this privileged EXEC command:

archive tar /xtract *source-url* **flash:***/file-url* [*dir/file...*]

For *source-url*, specify the source URL alias for the local file system. These options are supported:

- For the local flash file system, the syntax is
  flash:
- For the FTP, the syntax is
  ftp:[[*//username*[**:***password*]*@location*]*/directory*]*/tar-filename*.**tar**
- For the RCP, the syntax is
  rcp:[[*//username@location*]*/directory*]*/tar-filename*.**tar**
- For the TFTP, the syntax is
  tftp:[[*//location*]*/directory*]*/tar-filename*.**tar**

The *tar-filename*.**tar** is the tar file from which to extract files.

For **flash:***/file-url* [*dir/file...*], specify the location on the local flash file system into which the tar file is extracted. Use the *dir/file...* option to specify an optional list of files or directories within the tar file to be extracted. If none are specified, all files and directories are extracted.

This example shows how to extract the contents of a tar file located on the TFTP server at 172.20.10.30. This command extracts just the *new-configs* directory into the root directory on the local flash file system. The remaining files in the *saved.tar* file are ignored.

```
Switch# archive tar /xtract tftp://172.20.10.30/saved.tar flash:/new-configs
```

Displaying the Contents of a File

To display the contents of any readable file, including a file on a remote file system, use the **more** [**/ascii** | **/binary** | **/ebcdic**] *file-url* privileged EXEC command:.

This example shows how to display the contents of a configuration file on a TFTP server:

```
Switch# more tftp://serverA/hampton/savedconfig
!
! Saved configuration on server
!
version 11.3
service timestamps log datetime localtime
service linenumber
service udp-small-servers
service pt-vty-logging
!
<output truncated>
```

Working with Configuration Files

This section describes how to create, load, and maintain configuration files.

Configuration files contain commands entered to customize the function of the Cisco IOS software. A way to create a basic configuration file is to use the **setup** program or to enter the **setup** privileged EXEC command. For more information, see Chapter 4, "Performing Switch Setup Configuration."

You can copy (*download*) configuration files from a TFTP, FTP, or RCP server to the running configuration or startup configuration of the switch. You might want to perform this for one of these reasons:

- To restore a backed-up configuration file.
- To use the configuration file for another switch. For example, you might add another switch to your network and want it to have a configuration similar to the original switch. By copying the file to the new switch, you can change the relevant parts rather than recreating the whole file.
- To load the same configuration commands on all the switches in your network so that all the switches have similar configurations.

You can copy (*upload*) configuration files from the switch to a file server by using TFTP, FTP, or RCP. You might perform this task to back up a current configuration file to a server before changing its contents so that you can later restore the original configuration file from the server.

The protocol you use depends on which type of server you are using. The FTP and RCP transport mechanisms provide faster performance and more reliable delivery of data than TFTP. These improvements are possible because FTP and RCP are built on and use the TCP/IP stack, which is connection-oriented.

Guidelines for Creating and Using Configuration Files

Creating configuration files can aid in your switch configuration. Configuration files can contain some or all of the commands needed to configure one or more switches. For example, you might want to download the same configuration file to several switches that have the same hardware configuration.

Use these guidelines when creating a configuration file:

- We recommend that you connect through the console port for the initial configuration of the switch. If you are accessing the switch through a network connection instead of through a direct connection to the console port, keep in mind that some configuration changes (such as changing the switch IP address or disabling ports) can cause a loss of connectivity to the switch.
- If no password has been set on the switch, we recommend that you set one by using the **enable secret** *secret-password* global configuration command.

Note The **copy** {**ftp: | rcp: | tftp:**} **system:running-config** privileged EXEC command loads the configuration files on the switch as if you were entering the commands at the command line. The switch does not erase the existing running configuration before adding the commands. If a command in the copied configuration file replaces a command in the existing configuration file, the existing command is erased. For example, if the copied configuration file contains a different IP address in a particular command than the existing configuration, the IP address in the copied configuration is used. However, some commands in the existing configuration might not be replaced or negated. In this case, the resulting configuration file is a mixture of the existing configuration file and the copied configuration file, with the copied configuration file having precedence.

To restore a configuration file to an exact copy of a file stored on a server, copy the configuration file directly to the startup configuration (by using the **copy** {**ftp: | rcp: | tftp:**} **nvram:startup-config** privileged EXEC command), and reload the switch.

Configuration File Types and Location

Startup configuration files are used during system startup to configure the software. Running configuration files contain the current configuration of the software. The two configuration files can be different. For example, you might want to change the configuration for a short time period rather than permanently. In this case, you would change the running configuration but not save the configuration by using the **copy running-config startup-config** privileged EXEC command.

The running configuration is saved in DRAM; the startup configuration is stored in the NVRAM section of flash memory.

Creating a Configuration File By Using a Text Editor

When creating a configuration file, you must list commands logically so that the system can respond appropriately. This is one method of creating a configuration file:

Step 1 Copy an existing configuration from a switch to a server.

For more information, see the "Downloading the Configuration File By Using TFTP" section on page A-11, the "Downloading a Configuration File By Using FTP" section on page A-14, or the "Downloading a Configuration File By Using RCP" section on page A-17.

Step 2 Open the configuration file in a text editor, such as vi or emacs on UNIX or Notepad on a PC.

Step 3 Extract the portion of the configuration file with the desired commands, and save it in a new file.

Step 4 Copy the configuration file to the appropriate server location. For example, copy the file to the TFTP directory on the workstation (usually /tftpboot on a UNIX workstation).

Step 5 Make sure the permissions on the file are set to world-read.

Copying Configuration Files By Using TFTP

You can configure the switch by using configuration files you create, download from another switch, or download from a TFTP server. You can copy (upload) configuration files to a TFTP server for storage.

Preparing to Download or Upload a Configuration File By Using TFTP

Before you begin downloading or uploading a configuration file by using TFTP, do these tasks:

- Ensure that the workstation acting as the TFTP server is properly configured. On a Sun workstation, make sure that the /etc/inetd.conf file contains this line:

  ```
  tftp dgram udp wait root /usr/etc/in.tftpd in.tftpd -p -s /tftpboot
  ```

  Make sure that the /etc/services file contains this line:

  ```
  tftp 69/udp
  ```

  Note You must restart the inetd daemon after modifying the /etc/inetd.conf and /etc/services files. To restart the daemon, either stop the inetd process and restart it, or enter a **fastboot** command (on the SunOS 4.x) or a **reboot** command (on Solaris 2.x or SunOS 5.x). For more information on the TFTP daemon, see the documentation for your workstation.

- Ensure that the switch has a route to the TFTP server. The switch and the TFTP server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the TFTP server by using the **ping** command.
- Ensure that the configuration file to be downloaded is in the correct directory on the TFTP server (usually /tftpboot on a UNIX workstation).
- For download operations, ensure that the permissions on the file are set correctly. The permission on the file should be world-read.
- Before uploading the configuration file, you might need to create an empty file on the TFTP server. To create an empty file, enter the **touch** *filename* command, where *filename* is the name of the file you will use when uploading it to the server.
- During upload operations, if you are overwriting an existing file (including an empty file, if you had to create one) on the server, ensure that the permissions on the file are set correctly. Permissions on the file should be world-write.

Downloading the Configuration File By Using TFTP

To configure the switch by using a configuration file downloaded from a TFTP server, follow these steps:

Step 1 Copy the configuration file to the appropriate TFTP directory on the workstation.

Step 2 Verify that the TFTP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using TFTP" section on page A-11.

Step 3 Log into the switch through the console port or a Telnet session.

Step 4 Download the configuration file from the TFTP server to configure the switch.

Specify the IP address or hostname of the TFTP server and the name of the file to download.

Use one of these privileged EXEC commands:

- **copy tftp:**[[[*//location*]*/directory*]*/filename*] **system:running-config**
- **copy tftp:**[[[*//location*]*/directory*]*/filename*] **nvram:startup-config**

The configuration file downloads, and the commands are executed as the file is parsed line-by-line.

This example shows how to configure the software from the file *tokyo-confg* at IP address 172.16.2.155:

```
Switch# copy tftp://172.16.2.155/tokyo-confg system:running-config
Configure using tokyo-confg from 172.16.2.155? [confirm] y
Booting tokyo-confg from 172.16.2.155:!!! [OK - 874/16000 bytes]
```

Uploading the Configuration File By Using TFTP

To upload a configuration file from a switch to a TFTP server for storage, follow these steps:

Step 1 Verify that the TFTP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using TFTP" section on page A-11.

Step 2 Log into the switch through the console port or a Telnet session.

Step 3 Upload the switch configuration to the TFTP server. Specify the IP address or hostname of the TFTP server and the destination filename.

Use one of these privileged EXEC commands:

- **copy system:running-config tftp:**[[[*//location*]*/directory*]*/filename*]
- **copy nvram:startup-config tftp:**[[[*//location*]*/directory*]*/filename*]

The file is uploaded to the TFTP server.

This example shows how to upload a configuration file from a switch to a TFTP server:

```
Switch# copy system:running-config tftp://172.16.2.155/tokyo-confg
Write file tokyo-confg on host 172.16.2.155? [confirm] y
#
Writing tokyo-confg!!! [OK]
```

Copying Configuration Files By Using FTP

You can copy configuration files to or from an FTP server.

The FTP protocol requires a client to send a remote username and password on each FTP request to a server. When you copy a configuration file from the switch to a server by using FTP, the Cisco IOS software sends the first valid username in this list:

- The username specified in the **copy** command if a username is specified.
- The username set by the **ip ftp username** *username* global configuration command if the command is configured.
- Anonymous.

The switch sends the first valid password in this list:

- The password specified in the **copy** command if a password is specified.
- The password set by the **ip ftp password** *password* global configuration command if the command is configured.
- The switch forms a password named *username@switchname.domain*. The variable *username* is the username associated with the current session, *switchname* is the configured hostname, and *domain* is the domain of the switch.

The username and password must be associated with an account on the FTP server. If you are writing to the server, the FTP server must be properly configured to accept your FTP write request.

Use the **ip ftp username** and **ip ftp password** commands to specify a username and password for all copies. Include the username in the **copy** command if you want to specify only a username for that copy operation.

If the server has a directory structure, the configuration file is written to or copied from the directory associated with the username on the server. For example, if the configuration file resides in the home directory of a user on the server, specify that user's name as the remote username.

For more information, see the documentation for your FTP server.

Preparing to Download or Upload a Configuration File By Using FTP

Before you begin downloading or uploading a configuration file by using FTP, do these tasks:

- Ensure that the switch has a route to the FTP server. The switch and the FTP server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the FTP server by using the **ping** command.
- If you are accessing the switch through the console or a Telnet session and you do not have a valid username, make sure that the current FTP username is the one that you want to use for the FTP download. You can enter the **show users** privileged EXEC command to view the valid username. If you do not want to use this username, create a new FTP username by using the **ip ftp username** *username* global configuration command during all copy operations. The new username is stored in NVRAM. If you are accessing the switch through a Telnet session and you have a valid username, this username is used, and you do not need to set the FTP username. Include the username in the **copy** command if you want to specify a username for only that copy operation.
- When you upload a configuration file to the FTP server, it must be properly configured to accept the write request from the user on the switch.

For more information, see the documentation for your FTP server.

Downloading a Configuration File By Using FTP

Beginning in privileged EXEC mode, follow these steps to download a configuration file by using FTP:

	Command	Purpose
Step 1	Verify that the FTP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using FTP" section on page A-13.	
Step 2	Log into the switch through the console port or a Telnet session.	
Step 3	**configure terminal**	Enters global configuration mode on the switch. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
Step 4	**ip ftp username** *username*	(Optional) Changes the default remote username.
Step 5	**ip ftp password** *password*	(Optional) Changes the default password.
Step 6	**end**	Returns to privileged EXEC mode.
Step 7	**copy ftp:**[[[//[*username*[**:***password*]@]*location*]*/directory*] */filename*] **system:running-config** or **copy ftp:**[[[//[*username*[**:***password*]@]*location*]*/directory*] */filename*] **nvram:startup-config**	Using FTP, copies the configuration file from a network server to the running configuration or to the startup configuration file.

This example shows how to copy a configuration file named *host1-confg* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 and to load and run those commands on the switch:

```
Switch# copy ftp://netadmin1:mypass@172.16.101.101/host1-confg system:running-config
Configure using host1-confg from 172.16.101.101? [confirm]
Connected to 172.16.101.101
Loading 1112 byte file host1-confg:![OK]
Switch#
%SYS-5-CONFIG: Configured from host1-config by ftp from 172.16.101.101
```

This example shows how to specify a remote username of *netadmin1*. The software copies the configuration file *host2-confg* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 to the switch startup configuration.

```
Switch# configure terminal
Switch(config)# ip ftp username netadmin1
Switch(config)# ip ftp password mypass
Switch(config)# end
Switch# copy ftp: nvram:startup-config
Address of remote host [255.255.255.255]? 172.16.101.101
Name of configuration file[rtr2-confg]? host2-confg
Configure using host2-confg from 172.16.101.101?[confirm]
Connected to 172.16.101.101
Loading 1112 byte file host2-confg:![OK]
[OK]
Switch#
```

```
%SYS-5-CONFIG_NV:Non-volatile store configured from host2-config by ftp from
172.16.101.101
```

Uploading a Configuration File By Using FTP

Beginning in privileged EXEC mode, follow these steps to upload a configuration file by using FTP:

	Command	Purpose
Step 1	Verify that the FTP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using FTP" section on page A-13.	
Step 2	Log into the switch through the console port or a Telnet session.	
Step 3	**configure terminal**	Enters global configuration mode. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
Step 4	**ip ftp username** *username*	(Optional) Changes the default remote username.
Step 5	**ip ftp password** *password*	(Optional) Changes the default password.
Step 6	**end**	Returns to privileged EXEC mode.
Step 7	**copy system:running-config ftp:**[[[**//**[*username*[**:***password*]**@**]*location*]**/***directory*]**/***filename*] or **copy nvram:startup-config ftp:**[[[**//**[*username*[**:***password*]**@**]*location*]**/***directory*]**/***filename*]	Using FTP, copies the switch running or startup configuration file to the specified location.

This example shows how to copy the running configuration file named *switch2-confg* to the *netadmin1* directory on the remote host with an IP address of 172.16.101.101:

```
Switch# copy system:running-config ftp://netadmin1:mypass@172.16.101.101/switch2-confg
Write file switch2-confg on host 172.16.101.101?[confirm]
Building configuration...[OK]
Connected to 172.16.101.101
Switch#
```

This example shows how to store a startup configuration file on a server by using FTP to copy the file:

```
Switch# configure terminal
Switch(config)# ip ftp username netadmin2
Switch(config)# ip ftp password mypass
Switch(config)# end
Switch# copy nvram:startup-config ftp:
Remote host[]? 172.16.101.101
Name of configuration file to write [switch2-confg]?
Write file switch2-confg on host 172.16.101.101?[confirm]
![OK]
```

Copying Configuration Files By Using RCP

The RCP provides another method of downloading, uploading, and copying configuration files between remote hosts and the switch. Unlike TFTP, which uses User Datagram Protocol (UDP), a connectionless protocol, RCP uses TCP, which is connection-oriented.

To use RCP to copy files, the server from or to which you will be copying files must support RCP. The RCP copy commands rely on the rsh server (or daemon) on the remote system. To copy files by using RCP, you do not need to create a server for file distribution as you do with TFTP. You only need to have access to a server that supports the remote shell (rsh). (Most UNIX systems support rsh.) Because you are copying a file from one place to another, you must have read permission on the source file and write permission on the destination file. If the destination file does not exist, RCP creates it for you.

The RCP requires a client to send a remote username with each RCP request to a server. When you copy a configuration file from the switch to a server, the Cisco IOS software sends the first valid username in this list:

- The username specified in the **copy** command if a username is specified.
- The username set by the **ip rcmd remote-username** *username* global configuration command if the command is configured.
- The remote username associated with the current TTY (terminal) process. For example, if the user is connected to the router through Telnet and was authenticated through the **username** command, the switch software sends the Telnet username as the remote username.
- The switch hostname.

For a successful RCP copy request, you must define an account on the network server for the remote username. If the server has a directory structure, the configuration file is written to or copied from the directory associated with the remote username on the server. For example, if the configuration file is in the home directory of a user on the server, specify that user's name as the remote username.

Preparing to Download or Upload a Configuration File By Using RCP

Before you begin downloading or uploading a configuration file by using RCP, do these tasks:

- Ensure that the workstation acting as the RCP server supports the remote shell (rsh).
- Ensure that the switch has a route to the RCP server. The switch and the server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the RCP server by using the **ping** command.
- If you are accessing the switch through the console or a Telnet session and you do not have a valid username, make sure that the current RCP username is the one that you want to use for the RCP download. You can enter the **show users** privileged EXEC command to view the valid username. If you do not want to use this username, create a new RCP username by using the **ip rcmd remote-username** *username* global configuration command to be used during all copy operations. The new username is stored in NVRAM. If you are accessing the switch through a Telnet session and you have a valid username, this username is used, and you do not need to set the RCP username. Include the username in the **copy** command if you want to specify a username for only that copy operation.
- When you upload a file to the RCP server, it must be properly configured to accept the RCP write request from the user on the switch. For UNIX systems, you must add an entry to the .rhosts file for the remote user on the RCP server. For example, suppose that the switch contains these configuration lines:

```
hostname Switch1
```

```
ip rcmd remote-username User0
```

If the switch IP address translates to *Switch1.company.com*, the .rhosts file for User0 on the RCP server should contain this line:

```
Switch1.company.com Switch1
```

For more information, see the documentation for your RCP server.

Downloading a Configuration File By Using RCP

Beginning in privileged EXEC mode, follow these steps to download a configuration file by using RCP:

	Command	Purpose
Step 1	Verify that the RCP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using RCP" section on page A-16.	
Step 2	Log into the switch through the console port or a Telnet session.	
Step 3	**configure terminal**	Enters global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
Step 4	**ip rcmd remote-username** *username*	(Optional) Specifesthe remote username.
Step 5	**end**	Returns to privileged EXEC mode.
Step 6	**copy rcp:**[[[**//**[*username*@]*location*]**/***directory*]**/***filename*] **system:running-config** or **copy rcp:**[[[**//**[*username*@]*location*]**/***directory*]**/***filename*] **nvram:startup-config**	Using RCP, copies the configuration file from a network server to the running configuration or to the startup configuration file.

This example shows how to copy a configuration file named *host1-confg* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 and load and run those commands on the switch:

```
Switch# copy rcp://netadmin1@172.16.101.101/host1-confg system:running-config
Configure using host1-confg from 172.16.101.101? [confirm]
Connected to 172.16.101.101
Loading 1112 byte file host1-confg:![OK]
Switch#
%SYS-5-CONFIG: Configured from host1-config by rcp from 172.16.101.101
```

This example shows how to specify a remote username of *netadmin1*. Then it copies the configuration file *host2-confg* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 to the startup configuration:

```
Switch# configure terminal
Switch(config)# ip rcmd remote-username netadmin1
Switch(config)# end
Switch# copy rcp: nvram:startup-config
```

```
Address of remote host [255.255.255.255]? 172.16.101.101
Name of configuration file[rtr2-confg]? host2-confg
Configure using host2-confg from 172.16.101.101?[confirm]
Connected to 172.16.101.101
Loading 1112 byte file host2-confg:![OK]
[OK]
Switch#
%SYS-5-CONFIG_NV:Non-volatile store configured from host2-config by rcp from
172.16.101.101
```

Uploading a Configuration File By Using RCP

Beginning in privileged EXEC mode, follow these steps to upload a configuration file by using RCP:

	Command	Purpose
Step 1	Verify that the RCP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using RCP" section on page A-16.	
Step 2	Log into the switch through the console port or a Telnet session.	
Step 3	**configure terminal**	Enters global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
Step 4	**ip rcmd remote-username** *username*	(Optional) Specifies the remote username.
Step 5	**end**	Returns to privileged EXEC mode.
Step 6	**copy system:running-config rcp:**[[[**//**[*username*@]*location*]**/***directory*]**/***filename*] or **copy nvram:startup-config rcp:**[[[**//**[*username*@]*location*]**/***directory*]**/***filename*]	Using RCP, copies the configuration file from a switch running or startup configuration file to a network server.

This example shows how to copy the running configuration file named *switch2-confg* to the *netadmin1* directory on the remote host with an IP address of 172.16.101.101:

```
Switch# copy system:running-config rcp://netadmin1@172.16.101.101/switch2-confg
Write file switch-confg on host 172.16.101.101?[confirm]
Building configuration...[OK]
Connected to 172.16.101.101
Switch#
```

This example shows how to store a startup configuration file on a server:

```
Switch# configure terminal
Switch(config)# ip rcmd remote-username netadmin2
Switch(config)# end
Switch# copy nvram:startup-config rcp:
Remote host[]? 172.16.101.101
Name of configuration file to write [switch2-confg]?
Write file switch2-confg on host 172.16.101.101?[confirm]
![OK]
```

Clearing Configuration Information

You can clear the configuration information from the startup configuration. If you reboot the switch with no startup configuration, the switch enters the setup program so that you can reconfigure the switch with all new settings.

Clearing the Startup Configuration File

To clear the contents of your startup configuration, use the **erase nvram:** or the **erase startup-config** privileged EXEC command.

Caution You cannot restore the startup configuration file after it has been deleted.

Deleting a Stored Configuration File

To delete a saved configuration from flash memory, use the **delete flash:***filename* privileged EXEC command. Depending on the setting of the **file prompt** global configuration command, you might be prompted for confirmation before you delete a file. By default, the switch prompts for confirmation on destructive file operations. For more information about the **file prompt** command, see the *Cisco IOS Command Reference for Release 12.2*.

Caution You cannot restore a file after it has been deleted.

Replacing and Rolling Back Configurations

The configuration replacement and rollback feature replaces the running configuration with any saved Cisco IOS configuration file. You can use the rollback function to roll back to a previous configuration.

Understanding Configuration Replacement and Rollback

Archiving a Configuration

The configuration archive provides a mechanism to store, organize, and manage an archive of configuration files. The **configure replace** privileged EXEC command increases the configuration rollback capability. As an alternative, you can save copies of the running configuration by using the **copy running-config** *destination-url* privileged EXEC command, storing the replacement file either locally or remotely. However, this method lacks any automated file management. The configuration replacement and rollback feature can automatically save copies of the running configuration to the configuration archive.

You use the **archive config** privileged EXEC command to save configurations in the configuration archive by using a standard location and filename prefix that is automatically appended with an incremental version number (and optional timestamp) as each consecutive file is saved. You can specify how many versions of the running configuration are kept in the archive. After the maximum number of files are saved, the oldest file is automatically deleted when the next, most recent file is saved. The **show archive** privileged EXEC command displays information for all the configuration files saved in the configuration archive.

The Cisco IOS configuration archive, in which the configuration files are stored and available for use with the **configure replace** command, is in any of these file systems: FTP, HTTP, RCP, TFTP.

Replacing a Configuration

The **configure replace** privileged EXEC command replaces the running configuration with any saved configuration file. When you enter the **configure replace** command, the running configuration is compared with the specified replacement configuration, and a set of configuration differences is generated. The resulting differences are used to replace the configuration. The configuration replacement operation is usually completed in no more than three passes. To prevent looping behavior no more than five passes are performed.

You can use the **copy** *source-url* **running-config** privileged EXEC command to copy a stored configuration file to the running configuration. When using this command as an alternative to the **configure replace** *target-url* privileged EXEC command, note these major differences:

- The **copy** *source-url* **running-config** command is a merge operation and preserves all the commands from both the source file and the running configuration. This command does not remove commands from the running configuration that are not present in the source file. In contrast, the **configure replace** *target-url* command removes commands from the running configuration that are not present in the replacement file and adds commands to the running configuration that are not present.
- You can use a partial configuration file as the source file for the **copy** *source-url* **running-config** command. You must use a complete configuration file as the replacement file for the **configure replace** *target-url* command.

Rolling Back a Configuration

You can also use the **configure replace** command to roll back changes that were made since the previous configuration was saved. Instead of basing the rollback operation on a specific set of changes that were applied, the configuration rollback capability reverts to a specific configuration based on a saved configuration file.

If you want the configuration rollback capability, you must first save the running configuration before making any configuration changes. Then, after entering configuration changes, you can use that saved configuration file to roll back the changes by using the **configure replace** *target-url* command.

You can specify any saved configuration file as the rollback configuration. You are not limited to a fixed number of rollbacks, as is the case in some rollback models.

Configuration Guidelines

Follow these guidelines when configuring and performing configuration replacement and rollback:

- Make sure that the switch has free memory larger than the combined size of the two configuration files (the running configuration and the saved replacement configuration). Otherwise, the configuration replacement operation fails.
- Make sure that the switch also has sufficient free memory to execute the configuration replacement or rollback configuration commands.
- Certain configuration commands, such as those pertaining to physical components of a networking device (for example, physical interfaces), cannot be added or removed from the running configuration.
  - A configuration replacement operation cannot remove the **interface** *interface-id* command line from the running configuration if that interface is physically present on the device.

  - The **interface** *interface-id* command line cannot be added to the running configuration if no such interface is physically present on the device.

- When using the **configure replace** command, you must specify a saved configuration as the replacement configuration file for the running configuration. The replacement file must be a complete configuration generated by a Cisco IOS device (for example, a configuration generated by the **copy running-config** *destination-url* command).

Note If you generate the replacement configuration file externally, it must comply with the format of files generated by Cisco IOS devices.

Configuring the Configuration Archive

Using the **configure replace** command with the configuration archive and with the **archive config** command is optional but offers significant benefit for configuration rollback scenarios. Before using the **archive config command**, you must first configure the configuration archive. Starting in privileged EXEC mode, follow these steps to configure the configuration archive:

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**archive**	Enters archive configuration mode.
Step 3	**path** *url*	Specifies the location and filename prefix for the files in the configuration archive.
Step 4	**maximum** *number*	(Optional) Sets the maximum number of archive files of the running configuration to be saved in the configuration archive. *number*—Maximum files of the running configuration file in the configuration archive. Valid values are from 1 to 14. The default is 10. **Note** Before using this command, you must first enter the **path** archive configuration command to specify the location and filename prefix for the files in the configuration archive.
Step 5	**time-period** *minutes*	(Optional) Sets the time increment for automatically saving an archive file of the running configuration in the configuration archive. *minutes*—Specifies how often, in minutes, to automatically save an archive file of the running configuration in the configuration archive.
Step 6	**end**	Returns to privileged EXEC mode.
Step 7	**show running-config**	Verifies the configuration.
Step 8	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Performing a Configuration Replacement or Rollback Operation

Starting in privileged EXEC mode, follow these steps to replace the running configuration file with a saved configuration file:

	Command	Purpose
Step 1	**archive config**	(Optional) Saves the running configuration file to the configuration archive. **Note** Enter the **path** archive configuration command before using this command.
Step 2	**configure terminal**	Enters global configuration mode.
Step 3		Makes necessary changes to the running configuration.
Step 4	**exit**	Returns to privileged EXEC mode.
Step 5	**configure replace** *target-url* [**list**] [**force**] [**time** *seconds*] [**nolock**]	Replaces the running configuration file with a saved configuration file. *target-url*—URL (accessible by the file system) of the saved configuration file that is to replace the running configuration, such as the configuration file created in Step 2 by using the **archive config** privileged EXEC command. **list**—Displays a list of the command entries applied by the software parser during each pass of the configuration replacement operation. The total number of passes also appears. **force**— Replaces the running configuration file with the specified saved configuration file without prompting you for confirmation. **time** *seconds*—Specifies the time (in seconds) within which you must enter the **configure confirm** command to confirm replacement of the running configuration file. If you do not enter the **configure confirm** command within the specified time limit, the configuration replacement operation is automatically stopped. (In other words, the running configuration file is restored to the configuration that existed before you entered the **configure replace** command). **Note** You must first enable the configuration archive before you can use the **time** *seconds* command line option. **nolock**—Disables the locking of the running configuration file that prevents other users from changing the running configuration during a configuration replacement operation.
Step 6	**configure confirm**	(Optional) Confirms replacement of the running configuration with a saved configuration file. **Note** Use this command only if the **time** *seconds* keyword and argument of the **configure replace** command are specified.
Step 7	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

Working with Software Images

This section describes how to archive (download and upload) software image files, which contain the system software, the Cisco IOS code, and the embedded Device Manager software.

Note Instead of using the **copy** privileged EXEC command or the **archive tar** privileged EXEC command, we recommend using the **archive download-sw** and **archive upload-sw** privileged EXEC commands to download and upload software image files.

You can download a switch image file from a TFTP, FTP, or RCP server to upgrade the switch software. If you do not have access to a TFTP server, you can download a software image file directly to your PC or workstation by using a web browser (HTTP) and then by using Device Manager or Cisco Network Assistant to upgrade your switch. For information about upgrading your switch by using a TFTP server or a web browser (HTTP), see the release notes.

You can replace the current image with the new one or keep the current image in flash memory after a download.

You upload a switch image file to a TFTP, FTP, or RCP server for backup purposes. You can use this uploaded image for future downloads to the same switch or to another of the same type.

The protocol that you use depends on which type of server you are using. The FTP and RCP transport mechanisms provide faster performance and more reliable delivery of data than TFTP. These improvements are possible because FTP and RCP are built on and use the TCP/IP stack, which is connection-oriented.

Note For a list of software images and the supported upgrade paths, see the release notes.

Image Location on the Switch

The Cisco IOS image is stored as a *.bin* file in a directory that shows the version number. A subdirectory contains the files needed for web management. The image is stored on the system board flash memory (flash:).

You can use the **show version** privileged EXEC command to see the software version that is currently running on your switch. In the display, check the line that begins with `System image file is...`. It shows the directory name in flash memory where the image is stored.

You can also use the **dir** *filesystem***:** privileged EXEC command to see the directory names of other software images that might be stored in flash memory.The **archive download-sw /directory** privileged EXEC command allows you to specify a directory one time followed by a tar file or list of tar files to be downloaded instead of specifying complete paths with each tar file.

tar File Format of Images on a Server or Cisco.com

Software images located on a server or downloaded from Cisco.com are provided in a tar file format, which contains these files:

- An *info* file, which serves as a table of contents for the tar file
- One or more subdirectories containing other images and files, such as Cisco IOS images and web management files

This example shows some of the information contained in the info file. Table A-4 provides additional details about this information:

```
system_type:0x00000000:image-name
    image_family:xxxx
```

```
    stacking_number:x
    info_end:
version_suffix:xxxx
    version_directory:image-name
    image_system_type_id:0x00000000
    image_name:image-nameB.bin
    ios_image_file_size:6398464
    total_image_file_size:8133632
    image_feature:IP|LAYER_3|PLUS|MIN_DRAM_MEG=128
    image_family:xxxx
    stacking_number:x
   board_ids:0x401100c4 0x00000000 0x00000001 0x00000003 0x00000002 0x00008000 0x00008002
0x40110000
    info_end:
```

Note Disregard the stacking_number field. It does not apply to the switch.

Table A-4 info File Description

Field	Description
version_suffix	Specifies the Cisco IOS image version string suffix.
version_directory	Specifies the directory where the Cisco IOS image and the HTML subdirectory are installed.
image_name	Specifies the name of the Cisco IOS image within the tar file.
ios_image_file_size	Specifies the Cisco IOS image size in the tar file, which is an approximate measure of how much flash memory is required to hold just the Cisco IOS image.
total_image_file_size	Specifies the size of all the images (the Cisco IOS image and the web management files) in the tar file, which is an approximate measure of how much flash memory is required to hold them.
image_feature	Describes the core functionality of the image.
image_min_dram	Specifies the minimum amount of DRAM needed to run this image.
image_family	Describes the family of products on which the software can be installed.

Copying Image Files By Using TFTP

You can download a switch image from a TFTP server or upload the image from the switch to a TFTP server.

You download a switch image file from a server to upgrade the switch software. You can overwrite the current image with the new one or keep the current image after a download.

You upload a switch image file to a server for backup purposes; this uploaded image can be used for future downloads to the same or another switch of the same type.

Note Instead of using the **copy** privileged EXEC command or the **archive tar** privileged EXEC command, we recommend using the **archive download-sw** and **archive upload-sw** privileged EXEC commands to download and upload software image files.

Preparing to Download or Upload an Image File By Using TFTP

Before you begin downloading or uploading an image file by using TFTP, do these tasks:

- Ensure that the workstation acting as the TFTP server is properly configured. On a Sun workstation, make sure that the /etc/inetd.conf file contains this line:

  ```
  tftp dgram udp wait root /usr/etc/in.tftpd in.tftpd -p -s /tftpboot
  ```

  Make sure that the /etc/services file contains this line:

  ```
  tftp 69/udp
  ```

  Note You must restart the inetd daemon after modifying the /etc/inetd.conf and /etc/services files. To restart the daemon, either stop the inetd process and restart it, or enter a **fastboot** command (on the SunOS 4.x) or a **reboot** command (on Solaris 2.x or SunOS 5.x). For more information on the TFTP daemon, see the documentation for your workstation.

- Ensure that the switch has a route to the TFTP server. The switch and the TFTP server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the TFTP server by using the **ping** command.
- Ensure that the image to be downloaded is in the correct directory on the TFTP server (usually /tftpboot on a UNIX workstation).
- For download operations, ensure that the permissions on the file are set correctly. The permission on the file should be world-read.
- Before uploading the image file, you might need to create an empty file on the TFTP server. To create an empty file, enter the **touch** *filename* command, where *filename* is the name of the file you will use when uploading the image to the server.
- During upload operations, if you are overwriting an existing file (including an empty file, if you had to create one) on the server, ensure that the permissions on the file are set correctly. Permissions on the file should be world-write.

Downloading an Image File By Using TFTP

You can download a new image file and replace the current image or keep the current image.

Beginning in privileged EXEC mode, follow Steps 1 through 3 to download a new image from a TFTP server and overwrite the existing image. To keep the current image, go to Step 3.

	Command	Purpose
Step 1	Copy the image to the appropriate TFTP directory on the workstation. Make sure that the TFTP server is properly configured; see the "Preparing to Download or Upload an Image File By Using TFTP" section on page A-25.	
Step 2	Log into the switch through the console port or a Telnet session.	

	Command	Purpose
Step 3	**archive download-sw /overwrite /reload tftp:**[[*//location*]*/directory*]*/image-name*.**tar**	Downloads the image file from the TFTP server to the switch, and overwrite the current image. • The **/overwrite** option overwrites the software image in flash memory with the downloaded image. • The **/reload** option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For **//***location*, specify the IP address of the TFTP server. • For */directory/image-name*.**tar**, specify the directory (optional) and the image to download. Directory and image names are case sensitive.
Step 4	**archive download-sw /leave-old-sw /reload tftp:**[[*//location*]*/directory*]*/image-name*.**tar**	Downloads the image file from the TFTP server to the switch, and keep the current image. • The **/leave-old-sw** option keeps the old software version after a download. • The **/reload** option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For **//***location*, specify the IP address of the TFTP server. • For */directory/image-name*.**tar**, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it aborts the process and reports an error. If you specify the **/overwrite** option, the download algorithm removes the existing image on the flash device whether or not it is the same as the new one, downloads the new image, and then reloads the software.

Note If the flash device has sufficient space to hold two images and you want to overwrite one of these images with the same version, you must specify the **/overwrite** option.

If you specify the **/leave-old-sw**, the existing files are not removed. If there is not enough space to install the new image and keep the running image, the download process stops, and an error message is displayed.

The algorithm installs the downloaded image on the system board flash device (flash:). The image is placed into a new directory named with the software version string, and the BOOT environment variable is updated to point to the newly installed image.

If you kept the old image during the download process (you specified the **/leave-old-sw** keyword), you can remove it by entering the **delete /force /recursive** *filesystem***:/***file-url* privileged EXEC command. For *filesystem*, use **flash:** for the system board flash device. For *file-url*, enter the directory name of the old image. All the files in the directory and the directory are removed.

Caution For the download and upload algorithms to operate properly, do *not* rename image names.

Uploading an Image File By Using TFTP

You can upload an image from the switch to a TFTP server. You can later download this image to the switch or to another switch of the same type.

Use the upload feature only if the web management pages associated with the embedded Device Manager have been installed with the existing image.

Beginning in privileged EXEC mode, follow these steps to upload an image to a TFTP server:

	Command	Purpose
Step 1	Make sure the TFTP server is properly configured; see the "Preparing to Download or Upload an Image File By Using TFTP" section on page A-25.	
Step 2	Log into the switch through the console port or a Telnet session.	
Step 3	**archive upload-sw tftp:**[[*//location*]*/directory*]*/image-name*.**tar**	Uploads the currently running switch image to the TFTP server. • For *//location*, specify the IP address of the TFTP server. • For */directory/image-name*.**tar**, specify the directory (optional) and the name of the software image to be uploaded. Directory and image names are case sensitive. The *image-name*.**tar** is the name of the software image to be stored on the server.

The **archive upload-sw** privileged EXEC command builds an image file on the server by uploading these files in order: info, the Cisco IOS image, and the web management files. After these files are uploaded, the upload algorithm creates the tar file format.

Caution For the download and upload algorithms to operate properly, do *not* rename image names.

Copying Image Files By Using FTP

You can download a switch image from an FTP server or upload the image from the switch to an FTP server.

You download a switch image file from a server to upgrade the switch software. You can overwrite the current image with the new one or keep the current image after a download.

You upload a switch image file to a server for backup purposes. You can use this uploaded image for future downloads to the switch or another switch of the same type.

Note Instead of using the **copy** privileged EXEC command or the **archive tar** privileged EXEC command, we recommend using the **archive download-sw** and **archive upload-sw** privileged EXEC commands to download and upload software image files.

Preparing to Download or Upload an Image File By Using FTP

You can copy images files to or from an FTP server.

The FTP protocol requires a client to send a remote username and password on each FTP request to a server. When you copy an image file from the switch to a server by using FTP, the Cisco IOS software sends the first valid username in this list:

- The username specified in the **archive download-sw** or **archive upload-sw** privileged EXEC command if a username is specified.
- The username set by the **ip ftp username** *username* global configuration command if the command is configured.
- Anonymous.

The switch sends the first valid password in this list:

- The password specified in the **archive download-sw** or **archive upload-sw** privileged EXEC command if a password is specified.
- The password set by the **ip ftp password** *password* global configuration command if the command is configured.
- The switch forms a password named *username@switchname.domain*. The variable *username* is the username associated with the current session, *switchname* is the configured hostname, and *domain* is the domain of the switch.

The username and password must be associated with an account on the FTP server. If you are writing to the server, the FTP server must be properly configured to accept the FTP write request from you.

Use the **ip ftp username** and **ip ftp password** commands to specify a username and password for all copies. Include the username in the **archive download-sw** or **archive upload-sw** privileged EXEC command if you want to specify a username only for that operation.

If the server has a directory structure, the image file is written to or copied from the directory associated with the username on the server. For example, if the image file resides in the home directory of a user on the server, specify that user's name as the remote username.

Before you begin downloading or uploading an image file by using FTP, do these tasks:

- Ensure that the switch has a route to the FTP server. The switch and the FTP server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the FTP server by using the **ping** command.
- If you are accessing the switch through the console or a Telnet session and you do not have a valid username, make sure that the current FTP username is the one that you want to use for the FTP download. You can enter the **show users** privileged EXEC command to view the valid username. If you do not want to use this username, create a new FTP username by using the **ip ftp username** *username* global configuration command. This new name will be used during all archive operations. The new username is stored in NVRAM. If you are accessing the switch through a Telnet session and you have a valid username, this username is used, and you do not need to set the FTP username. Include the username in the **archive download-sw** or **archive upload-sw** privileged EXEC command if you want to specify a username for that operation only.
- When you upload an image file to the FTP server, it must be properly configured to accept the write request from the user on the switch.

For more information, see the documentation for your FTP server.

Downloading an Image File By Using FTP

You can download a new image file and overwrite the current image or keep the current image.

Beginning in privileged EXEC mode, follow Steps 1 through 7 to download a new image from an FTP server and overwrite the existing image. To keep the current image, go to Step 7.

	Command	Purpose
Step 1	Verify that the FTP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using FTP" section on page A-13.	
Step 2	Log into the switch through the console port or a Telnet session.	
Step 3	**configure terminal**	Enters global configuration mode. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
Step 4	**ip ftp username** *username*	(Optional) Changes the default remote username.
Step 5	**ip ftp password** *password*	(Optional) Changes the default password.
Step 6	**end**	Returns to privileged EXEC mode.
Step 7	**archive download-sw /overwrite /reload ftp:**[[*//username*[**:***password*]*@location*]*/directory*] */image-name***.tar**	Downloads the image file from the FTP server to the switch, and overwrite the current image. • The **/overwrite** option overwrites the software image in flash memory with the downloaded image. • The **/reload** option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For **//***username*[**:***password*], specify the username and password; these must be associated with an account on the FTP server. • For *@location*, specify the IP address of the FTP server. • For *directory/image-name***.tar**, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

	Command	Purpose
Step 8	**archive download-sw /leave-old-sw /reload ftp:**[[**//***username*[**:***password*]@*location*]**/***directory*] **/***image-name***.tar**	Downloads the image file from the FTP server to the switch, and keep the current image. • The **/leave-old-sw** option keeps the old software version after a download. • The **/reload** option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For **//***username*[**:***password*], specify the username and password. These must be associated with an account on the FTP server. • For @*location*, specify the IP address of the FTP server. • For *directory***/***image-name***.tar**, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it aborts the process and reports an error. If you specify the **/overwrite** option, the download algorithm removes the existing image on the flash device, whether or not it is the same as the new one, downloads the new image, and then reloads the software.

Note If the flash device has sufficient space to hold two images and you want to overwrite one of these images with the same version, you must specify the **/overwrite** option.

If you specify the **/leave-old-sw**, the existing files are not removed. If there is not enough space to install the new image and keep the running image, the download process stops, and an error message is displayed.

The algorithm installs the downloaded image onto the system board flash device (flash:). The image is placed into a new directory named with the software version string, and the BOOT environment variable is updated to point to the newly installed image.

If you kept the old image during the download process (you specified the **/leave-old-sw** keyword), you can remove it by entering the **delete /force /recursive** *filesystem***:/***file-url* privileged EXEC command. For *filesystem*, use **flash:** for the system board flash device. For *file-url*, enter the directory name of the old software image. All the files in the directory and the directory are removed.

Caution For the download and upload algorithms to operate properly, do *not* rename image names.

Uploading an Image File By Using FTP

You can upload an image from the switch to an FTP server. You can later download this image to the same switch or to another switch of the same type.

Use the upload feature only if the web management pages associated with the embedded Device Manager have been installed with the existing image.

Beginning in privileged EXEC mode, follow these steps to upload an image to an FTP server:

	Command	Purpose
Step 1	Verify that the FTP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using FTP" section on page A-13.	
Step 2	Log into the switch through the console port or a Telnet session.	
Step 3	**configure terminal**	Enters global configuration mode. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
Step 4	**ip ftp username** *username*	(Optional) Changes the default remote username.
Step 5	**ip ftp password** *password*	(Optional) Changes the default password.
Step 6	**end**	Returns to privileged EXEC mode.
Step 7	**archive upload-sw ftp:**[[**//**[*username*[**:***password*]@]*location*]*/directory*]**/***image-name***.tar**	Uploads the currently running switch image to the FTP server. • For **//***username***:***password*, specify the username and password. These must be associated with an account on the FTP server. • For @*location*, specify the IP address of the FTP server. • For **/***directory***/***image-name***.tar**, specify the directory (optional) and the name of the software image to be uploaded. Directory and image names are case sensitive. The *image-name***.tar** is the name of the software image to be stored on the server.

The **archive upload-sw** command builds an image file on the server by uploading these files in order: info, the Cisco IOS image, and the web management files. After these files are uploaded, the upload algorithm creates the tar file format.

Caution For the download and upload algorithms to operate properly, do *not* rename image names.

Copying Image Files By Using RCP

You can download a switch image from an RCP server or upload the image from the switch to an RCP server.

You download a switch image file from a server to upgrade the switch software. You can overwrite the current image with the new one or keep the current image after a download.

You upload a switch image file to a server for backup purposes. You can use this uploaded image for future downloads to the same switch or another of the same type.

Note Instead of using the **copy** privileged EXEC command or the **archive tar** privileged EXEC command, we recommend using the **archive download-sw** and **archive upload-sw** privileged EXEC commands to download and upload software image files.

Preparing to Download or Upload an Image File By Using RCP

RCP provides another method of downloading and uploading image files between remote hosts and the switch. Unlike TFTP, which uses User Datagram Protocol (UDP), a connectionless protocol, RCP uses TCP, which is connection-oriented.

To use RCP to copy files, the server from or to which you will be copying files must support RCP. The RCP copy commands rely on the rsh server (or daemon) on the remote system. To copy files by using RCP, you do not need to create a server for file distribution as you do with TFTP. You only need to have access to a server that supports the remote shell (rsh). (Most UNIX systems support rsh.) Because you are copying a file from one place to another, you must have read permission on the source file and write permission on the destination file. If the destination file does not exist, RCP creates it for you.

RCP requires a client to send a remote username on each RCP request to a server. When you copy an image from the switch to a server by using RCP, the Cisco IOS software sends the first valid username in this list:

- The username specified in the **archive download-sw** or **archive upload-sw** privileged EXEC command if a username is specified.
- The username set by the **ip rcmd remote-username** *username* global configuration command if the command is entered.
- The remote username associated with the current TTY (terminal) process. For example, if the user is connected to the router through Telnet and was authenticated through the **username** command, the switch software sends the Telnet username as the remote username.
- The switch hostname.

For the RCP copy request to execute successfully, an account must be defined on the network server for the remote username. If the server has a directory structure, the image file is written to or copied from the directory associated with the remote username on the server. For example, if the image file resides in the home directory of a user on the server, specify that user's name as the remote username.

Before you begin downloading or uploading an image file by using RCP, do these tasks:

- Ensure that the workstation acting as the RCP server supports the remote shell (rsh).
- Ensure that the switch has a route to the RCP server. The switch and the server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the RCP server by using the **ping** command.
- If you are accessing the switch through the console or a Telnet session and you do not have a valid username, make sure that the current RCP username is the one that you want to use for the RCP download. You can enter the **show users** privileged EXEC command to view the valid username. If you do not want to use this username, create a new RCP username by using the **ip rcmd remote-username** *username* global configuration command to be used during all archive operations. The new username is stored in NVRAM. If you are accessing the switch through a Telnet session and you have a valid username, this username is used, and there is no need to set the RCP username. Include the username in the **archive download-sw** or **archive upload-sw** privileged EXEC command if you want to specify a username only for that operation.

- When you upload an image to the RCP to the server, it must be properly configured to accept the RCP write request from the user on the switch. For UNIX systems, you must add an entry to the .rhosts file for the remote user on the RCP server.

  For example, suppose the switch contains these configuration lines:

  ```
  hostname Switch1
  ip rcmd remote-username User0
  ```

  If the switch IP address translates to *Switch1.company.com*, the .rhosts file for User0 on the RCP server should contain this line:

  ```
  Switch1.company.com Switch1
  ```

For more information, see the documentation for your RCP server.

Downloading an Image File By Using RCP

You can download a new image file and replace or keep the current image.

Beginning in privileged EXEC mode, follow Steps 1 through 6 to download a new image from an RCP server and overwrite the existing image. To keep the current image, go to Step 6.

	Command	Purpose
Step 1	Verify that the RCP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using RCP" section on page A-16.	
Step 2	Log into the switch through the console port or a Telnet session.	
Step 3	**configure terminal**	Enters global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
Step 4	**ip rcmd remote-username** *username*	(Optional) Specifies the remote username.
Step 5	**end**	Returns to privileged EXEC mode.
Step 6	**archive download-sw /overwrite /reload rcp:**[[[**//**[*username*@]*location*]**/***directory*]**/***image-name***.tar**]	Downloads the image file from the RCP server to the switch, and overwrite the current image. • The **/overwrite** option overwrites the software image in flash memory with the downloaded image. • The **/reload** option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For **//***username, s*pecify the username. For the RCP copy request to execute successfully, an account must be defined on the network server for the remote username. • For @*location*, specify the IP address of the RCP server. • For **/***directory***/***image-name***.tar**, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

	Command	Purpose
Step 7	**archive download-sw /leave-old-sw /reload rcp:**[[[**//**[*username*@]*location*]**/***directory*]**/***image-name***.tar**]	Downloads the image file from the RCP server to the switch, and keep the current image. • The **/leave-old-sw** option keeps the old software version after a download. • The **/reload** option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For **//***username, s*pecify the username. For the RCP copy request to execute, an account must be defined on the network server for the remote username. • For @*location*, specify the IP address of the RCP server. • **For /***directory*]**/***image-name***.tar**, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it aborts the process and reports an error. If you specify the **/overwrite** option, the download algorithm removes the existing image on the flash device whether or not it is the same as the new one, downloads the new image, and then reloads the software.

Note If the flash device has sufficient space to hold two images and you want to overwrite one of these images with the same version, you must specify the **/overwrite** option.

If you specify the **/leave-old-sw**, the existing files are not removed. If there is not enough room to install the new image an keep the running image, the download process stops, and an error message is displayed.

The algorithm installs the downloaded image onto the system board flash device (flash:). The image is placed into a new directory named with the software version string, and the BOOT environment variable is updated to point to the newly installed image.

If you kept the old software during the download process (you specified the **/leave-old-sw** keyword), you can remove it by entering the **delete /force /recursive** *filesystem***:/***file-url* privileged EXEC command. For *filesystem*, use **flash:** for the system board flash device. For *file-url*, enter the directory name of the old software image. All the files in the directory and the directory are removed.

Caution For the download and upload algorithms to operate properly, do *not* rename image names.

Uploading an Image File By Using RCP

You can upload an image from the switch to an RCP server. You can later download this image to the same switch or to another switch of the same type.

The upload feature should be used only if the web management pages associated with the embedded device manager have been installed with the existing image.

Beginning in privileged EXEC mode, follow these steps to upload an image to an RCP server:

	Command	Purpose
Step 1	Verify that the RCP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using RCP" section on page A-16.	
Step 2	Log into the switch through the console port or a Telnet session.	
Step 3	**configure terminal**	Enters global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
Step 4	**ip rcmd remote-username** *username*	(Optional) Specifies the remote username.
Step 5	**end**	Returns to privileged EXEC mode.
Step 6	**archive upload-sw rcp:**[[[**//**[*username*@]*location*]**/***directory*]**/***image-name***.tar**]	Uploads the currently running switch image to the RCP server. • For **//***username,* specify the username; for the RCP copy request to execute, an account must be defined on the network server for the remote username. • For @*location*, specify the IP address of the RCP server. • For **/***directory*]**/***image-name***.tar**, specify the directory (optional) and the name of the software image to be uploaded. Directory and image names are case sensitive. • The *image-name***.tar** is the name of software image to be stored on the server.

The **archive upload-sw** privileged EXEC command builds an image file on the server by uploading these files in order: info, the Cisco IOS image, and the web management files. After these files are uploaded, the upload algorithm creates the tar file format.

Caution For the download and upload algorithms to operate properly, do *not* rename image names.

INDEX

Numerics

802.1x Accounting
  Overview 13-31

A

AAA down policy, NAC Layer 2 IP validation 1-9
abbreviating commands 2-4
access-class command 37-17
access control entries
  See ACEs
access-denied response, VMPS 17-14
access groups
  Layer 2 37-17
accessing
  clusters, switch 6-11
  member switches 6-11
  switch clusters 6-11
access lists
  See ACLs
access ports
  defined 15-3
  in switch clusters 6-10
accounting
  with IEEE 802.1x 13-13
  with RADIUS 12-16, 12-37
  with TACACS+ 12-6, 12-8, 12-33
ACEs
  and QoS 38-13
  defined 37-2
  Ethernet 37-2
  IP 37-2
ACLs
  ACEs 37-2
  any keyword 37-14
  applying
    time ranges to 37-8, 37-16
    to an interface 37-9, 37-17
    to QoS 38-13
  classifying traffic for QoS 38-36
  comments in 37-8
  defined 37-1, 37-5
  examples of 38-36
  extended IPv4
    creating 37-6, 37-13
    matching criteria 37-5
  hardware and software handling 37-10
  host keyword 37-14
  IP
    creating 37-5
    fragments and QoS guidelines 38-5
    implicit deny 37-7, 37-8, 37-11
    implicit masks 37-11
    matching criteria 37-5
    undefined 37-10
  IPv4
    applying to interfaces 37-9, 37-17
    creating 37-5
    matching criteria 37-5
    named 37-7, 37-15
    numbers 37-5
    terminal lines, setting on 37-9, 37-17
    unsupported features 37-1
  logging messages 37-6
  MAC extended 37-11

matching 37-5, 37-9
named, IPv4 37-7
number per QoS class map 38-5
port 37-2
QoS 38-13, 38-36
resequencing entries 37-7
standard IPv4
creating 37-11
matching criteria 37-5
support for 1-7
support in hardware 37-10
types supported 37-2
unsupported features, IPv4 37-1
active link 24-4, 24-10, 24-11
active links 24-1
active traffic monitoring, IP SLAs 45-1
address aging time for VLANs 7-6
address aliasing 28-2
addresses
dynamic
accelerated aging 20-8
changing the aging time 7-6
default aging 20-8
defined 7-5
learning 7-5
IPv6 42-2
MAC, discovering 7-8
multicast
STP address management 20-8
static
adding and removing 7-6
defined 7-5
address resolution 7-8
Address Resolution Protocol
See ARP
administrative VLAN, REP 23-9
advertisements
CDP 32-1
LLDP 31-2
VTP 17-11, 18-3
age timer, REP 23-8
aggregatable global unicast addresses 42-3
aggregated ports
See EtherChannel
aggregate policers 38-46
aggregate policing 1-10
aging, accelerating 20-8
aging time
accelerated
for STP 20-8
MAC address table 7-13
alarm profiles
creating or modifying 3-8
alarms
displaying 3-9
power supply 3-2
temperature 3-2
allowed-VLAN list 17-12
ARP
defined 1-4, 7-8
table
address resolution 7-8
managing 7-8
attributes, RADIUS
vendor-proprietary 12-17, 12-38
vendor-specific 12-16
attribute-value pairs 13-10, 13-13, 13-19
authentication
local mode with AAA 12-20, 12-39
open1x 13-28
RADIUS
key 12-15
login 12-15, 12-36
TACACS+
defined 12-6
key 12-7, 12-30
login 12-7, 12-31

authentication compatibility with Catalyst 6500 switches 13-7
authentication failed VLAN
See restricted VLAN
authentication manager
CLI commands 13-8
overview 13-6
authoritative time source, described 7-2
authorization
with RADIUS 12-16, 12-37
with TACACS+ 12-6, 12-7, 12-33
authorized ports with IEEE 802.1x 13-9
auto enablement 13-29
automatic discovery
considerations
beyond a noncandidate device 6-8
brand new switches 6-10
connectivity 6-5
different VLANs 6-7
management VLANs 6-8
non-CDP-capable devices 6-7
noncluster-capable devices 6-7
routed ports 6-9
in switch clusters 6-5
See also CDP
automatic QoS
See QoS
auto-MDIX
described 15-10
autonegotiation
duplex mode 1-2
interface configuration guidelines 15-9
mismatches 46-1
autosensing, port speed 1-2
auxiliary VLAN
See voice VLAN
availability, features 1-5

B

BackboneFast
described 22-5
backup interfaces
See Flex Links
backup links 24-1
banners
configuring
login 7-13
message-of-the-day login 7-12
when displayed 7-4
Berkeley r-tools replacement 12-24
binding database
address, DHCP server
See DHCP, Cisco IOS server database
DHCP snooping
See DHCP snooping binding database
bindings
address, Cisco IOS DHCP server 25-6
DHCP snooping database 25-6
IP source guard 27-2
binding table, DHCP snooping
See DHCP snooping binding database
blocking packets 29-4, 29-11
booting
boot loader, function of 4-2
boot process 4-1
specific image 4-17
boot loader
accessing 4-10
described 4-2
environment variables 4-10
prompt 4-10
trap-door mechanism 4-2
BPDU
error-disabled state 22-2
filtering 22-3
RSTP format 21-11

BPDU filtering
described 22-3
support for 1-6
BPDU guard
described 22-2
support for 1-6
bridge protocol data unit
See BPDU
broadcast storm-control command 29-9
broadcast storms 29-1

C

cables, monitoring for unidirectional links 33-1
candidate switch
automatic discovery 6-5
defined 6-2
requirements 6-2
See also command switch, cluster standby group, and member switch
Catalyst 6500 switches
authentication compatibility 13-7
CA trustpoint
configuring 12-23
defined 12-23
CDP
and trusted boundary 38-27
automatic discovery in switch clusters 6-5
configuring 32-2
defined with LLDP 31-1
described 32-1
monitoring 32-3
overview 32-1
support for 1-4
transmission timer and holdtime, setting 32-2
updates 32-2
CGMP
as IGMP snooping learning method 28-7
joining multicast group 28-3
switch support of 1-2
CIP, enabling 10-2
CIP configuration 10-1
CipherSuites 12-24
Cisco 7960 IP Phone 19-1
Cisco Discovery Protocol
See CDP
Cisco Group Management Protocol
See CGMP
Cisco IOS DHCP server
See DHCP, Cisco IOS DHCP server
Cisco IOS File System
See IFS
Cisco IOS IP SLAs 45-2
Cisco Secure ACS
attribute-value pairs for downloadable ACLs 13-19
attribute-value pairs for redirect URL 13-19
Cisco Secure ACS configuration guide 13-48
CiscoWorks 2000 1-3, 36-5
CISP 13-29
CIST regional root
See MSTP
CIST root
See MSTP
civic location 31-3
class maps for QoS
configuring 38-38
described 38-13
displaying 38-56
class of service
See CoS
clearing interfaces 15-19
CLI
abbreviating commands 2-4
command modes 2-1
configuration logging 2-5
described 1-3
editing features
enabling and disabling 2-7

keystroke editing 2-7
wrapped lines 2-9
error messages 2-5
filtering command output 2-10
getting help 2-3
history
changing the buffer size 2-6
described 2-6
disabling 2-7
recalling commands 2-6
managing clusters 6-13
no and default forms of commands 2-4
Client Information Signalling Protocol
See CISP
client mode, VTP 18-3
clock
See system clock
clusters, switch
accessing 6-11
automatic discovery 6-5
benefits 1-2
compatibility 6-4
LRE profile considerations 6-13
managing
through CLI 6-13
through SNMP 6-14
planning considerations
automatic discovery 6-5
CLI 6-13
host names 6-11
IP addresses 6-11
LRE profiles 6-13
passwords 6-12
RADIUS 6-12
SNMP 6-12, 6-14
TACACS+ 6-12
See also candidate switch, command switch, cluster standby group, member switch, and standby command switch
cluster standby group
defined 6-3
requirements 6-2
CNS 1-4
Configuration Engine
configID, deviceID, hostname 5-4
configuration service 5-3
event service 5-3
embedded agents
described 5-5
enabling configuration agent 5-8
enabling event agent 5-7
management functions 1-3
CoA Request Commands 12-12
command-line interface
See CLI
command modes 2-1
commands
abbreviating 2-4
no and default 2-4
commands, setting privilege levels 12-29
command switch
configuration conflicts 46-9
defined 6-3
password privilege levels 6-13
recovery
from lost member connectivity 46-9
requirements 6-1
See also candidate switch, cluster standby group, member switch, and standby command switch
Common Industrial Protocol (CIP) 10-1
community strings
configuring 6-12, 36-6, 36-9
for cluster switches 36-4
in clusters 6-12
overview 36-4
SNMP 6-12
compatibility, feature 29-7
config.text 4-3

configurable leave timer, IGMP 28-5
configuration, initial
defaults 1-11
Express Setup 1-2
configuration changes, logging 35-9
configuration conflicts, recovering from lost member connectivity 46-9
configuration examples, network 1-14
configuration files
archiving A-19
clearing the startup configuration A-19
creating using a text editor A-10
default name 4-3
deleting a stored configuration A-19
described A-9
downloading
preparing A-11, A-13, A-16
reasons for A-9
using FTP A-14
using RCP A-17
using TFTP A-11
guidelines for creating and using A-9
guidelines for replacing and rolling back A-20
invalid combinations when copying A-6
obtaining with DHCP 4-9
password recovery disable considerations 12-3
replacing a running configuration A-19, A-20
rolling back a running configuration A-19, A-20
specifying the filename 4-15
types and location A-10
uploading
preparing A-11, A-13, A-16
reasons for A-9
using FTP A-15
using RCP A-18
using TFTP A-12
configuration guidelines
REP 23-7
configuration logger 35-9
configuration logging 2-5
configuration replacement A-19
configuration rollback A-19
configuration settings, saving 4-17
configure terminal command 15-13
configuring 802.1x user distribution 13-46
configuring small-frame arrival rate 29-10
config-vlan mode 2-2
conflicts, configuration 46-9
connections, secure remote 12-21
consistency checks in VTP Version 2 18-5
console port, connecting to 2-10
control protocol, IP SLAs 45-3
convergence
REP 23-4
corrupted software, recovery steps with Xmodem 46-7
CoS
override priority 19-4
trust priority 19-4
CoS input queue threshold map for QoS 38-22
CoS output queue threshold map for QoS 38-24
CoS-to-DSCP map for QoS 38-47
counters, clearing interface 15-19
CPU utilization, troubleshooting 46-6
crashinfo file 46-5
critical authentication, IEEE 802.1x 13-44
critical VLAN 13-22
cryptographic software image
Kerberos 12-17
SSH 12-1, 12-21
SSL 12-22
customjzeable web pages, web-based authentication 14-6

D

DACL
See downloadable ACL
daylight saving time 7-10
debugging

enabling all system diagnostics 46-12
enabling for a specific feature 46-12
redirecting error message output 46-13
default commands 2-4
default configuration
802.1x 13-30
auto-QoS 39-3
DHCP 25-7
DHCP option 82 25-8
DHCP snooping 25-8
DHCP snooping binding database 25-8
DNS 7-4
EtherChannel 40-10
Ethernet interfaces 15-8
Flex Links 24-5
IGMP filtering 28-12
IGMP snooping 28-6, 44-5
IGMP throttling 28-12
Layer 2 interfaces 15-15
LLDP 31-4
MAC address-table move update 24-5
MSTP 21-13
MVR 28-11
optional spanning-tree configuration 22-9
password and privilege level 12-2
PROFINET 9-4
PTP 8-2
RADIUS 12-10
REP 23-7
RSPAN 30-10
SNMP 36-8
SPAN 30-10
SSL 12-23
standard QoS 38-6
STP 20-11
system message logging 35-5
TACACS+ 12-7
UDLD 33-4
VLANs 17-6
VMPS 17-15
voice VLAN 19-3
VTP 18-9
default gateway 4-14, 4-15
default router preference
See DRP
default web-based authentication configuration
802.1X 14-10
deleting VLANs 17-7, 17-17
denial-of-service attack 29-1
designing your network, examples 1-14
destination addresses
in IPv4 ACLs 37-13
destination-IP address-based forwarding, EtherChannel 40-8
destination-MAC address forwarding, EtherChannel 40-8
detecting indirect link failures, STP 22-5
device discovery protocol 31-1, 32-1
device manager
benefits 1-2
described 1-2, 1-3
in-band management 1-5
DHCP
Cisco IOS server database
default configuration 25-8
described 25-6
enabling
relay agent 25-10
DHCP-based autoconfiguration
client request message exchange 4-5
configuring
client side 4-5
DNS Server 4-8
relay device 4-8
server side 4-7
TFTP server 4-7
lease options
for IP address information 4-7
for receiving the configuration file 4-7

relationship to BOOTP 4-5
relay support 1-4
support for 1-4
DHCP binding database
See DHCP snooping binding database
DHCP binding table
See DHCP snooping binding database
DHCP option 82
circuit ID suboption 25-4
configuration guidelines 25-8
default configuration 25-7
forwarding address, specifying 25-9, 25-10
helper address 25-9
overview 25-3
packet format, suboption
circuit ID 25-4
remote ID 25-4
remote ID suboption 25-4
DHCP server port-based address allocation
default configuration 25-10
described 25-9
enabling 25-13
DHCP server port-based address assignment
support for 1-4
DHCP snooping
accepting untrusted packets form edge switch 25-3, 25-11
binding database
See DHCP snooping binding database
configuration guidelines 25-8
default configuration 25-7
message exchange process 25-4
option 82 data insertion 25-3
trusted interface 25-2
untrusted interface 25-2
untrusted messages 25-2
DHCP snooping binding database
binding file
format 25-7
location 25-6
bindings 25-6
default configuration 25-7, 25-8
described 25-6
displaying
status and statistics 25-14
entry 25-6
DHCP snooping binding table
See DHCP snooping binding database
Differentiated Services architecture, QoS 38-2
Differentiated Services Code Point 38-2
directed unicast requests 1-4
directories
changing A-5
creating and removing A-5
displaying the working A-5
discovery, clusters
See automatic discovery
displaying switch alarms 3-9
DNS
and DHCP-based autoconfiguration 4-8
default configuration 7-4
in IPv6 42-3
overview 7-4
support for 1-4
domain names
DNS 7-4
VTP 18-10
Domain Name System
See DNS
downloadable ACL 13-18, 13-19, 13-48
downloading
configuration files
preparing A-11, A-13, A-16
reasons for A-9
using FTP A-14
using RCP A-17
using TFTP A-11
image files

deleting old image A-26
preparing A-25, A-28, A-32
reasons for A-23
using FTP A-29
using HTTP A-23
using RCP A-33
using TFTP A-25
using the device manager or Network Assistant A-23
DRP
configuring 42-8
described 42-4
IPv6 42-4
DSCP 1-10, 38-2
DSCP input queue threshold map for QoS 38-22
DSCP output queue threshold map for QoS 38-24
DSCP-to-CoS map for QoS 38-48
DSCP-to-DSCP-mutation map for QoS 38-30, 38-49
DSCP transparency 38-27, 38-34
DTP 1-6, 17-9
dual-action detection 40-5
dual IPv4 and IPv6 templates 11-3, 42-5
dual protocol stacks
IPv4 and IPv6 42-5
SDM templates supporting 42-5
dual-purpose uplinks
defined 15-4
LEDs 15-5
link selection 15-5
setting the type 15-15
dynamic access ports
characteristics 17-4
configuring 17-23
defined 15-3
dynamic addresses
See addresses
dynamic ARP inspection
ARP cache poisoning 26-2
ARP requests, described 26-1
ARP spoofing attack 26-2
configuration guidelines 26-5
configuring
ACLs for non-DHCP environments 26-7
in DHCP environments 26-6
log buffer 26-11
rate limit for incoming ARP packets 26-4, 26-9
described 26-1
DHCP snooping binding database 26-2
displaying
ARP ACLs 26-12
configuration and operating state 26-12
trust state and rate limit 26-12
error-disabled state for exceeding rate limit 26-4
function of 26-2
interface trust states 26-3
log buffer
configuring 26-11
logging of dropped packets, described 26-4
man-in-the middle attack, described 26-2
network security issues and interface trust states 26-3
priority of ARP ACLs and DHCP snooping entries 26-4
rate limiting of ARP packets
configuring 26-9
described 26-4
error-disabled state 26-4
validation checks, performing 26-10
dynamic auto trunking mode 17-10
dynamic desirable trunking mode 17-10
Dynamic Host Configuration Protocol
See DHCP-based autoconfiguration
dynamic port VLAN membership
described 17-15
reconfirming 17-16
troubleshooting 17-16
types of connections 17-23
Dynamic Trunking Protocol
See DTP

E

editing features
  enabling and disabling 2-7
  keystrokes used 2-7
  wrapped lines 2-9
ELIN location 31-3
enable password 12-3
enable secret password 12-3
enabling SNMP traps 3-9
encryption, CipherSuite 12-24
encryption for passwords 12-3, 12-27
environment variables, function of 4-11
error-disabled state, BPDU 22-2
error messages during command entry 2-5
EtherChannel
  automatic creation of 40-4, 40-6
  channel groups
    binding physical and logical interfaces 40-3
    numbering of 40-3
  configuration guidelines 40-10
  configuring
    Layer 2 interfaces 40-11
  default configuration 40-10
  forwarding methods 40-8, 40-14
  IEEE 802.3ad, described 40-6
  interaction
    with STP 40-10
    with VLANs 40-11
  LACP
    described 40-6
    hot-standby ports 40-7
    interaction with other features 40-7
    modes 40-6
  load balancing 40-8, 40-14
  logical interfaces, described 40-3
  PAgP
    aggregate-port learners 40-5
    compatibility with Catalyst 1900 40-5
    described 40-4
    interaction with other features 40-6
    interaction with virtual switches 40-5
    learn method and priority configuration 40-5, 40-14
    modes 40-4
    support for 1-2
    with dual-action detection 40-5
  port-channel interfaces
    described 40-3
  port groups 15-4
EtherChannel guard
  described 22-7
Ethernet VLANs
  adding 17-17
  defaults and ranges 17-7
  modifying 17-17
EUI 42-3
examples
  network configuration 1-14
expedite queue for QoS 38-56
Express Setup 1-2
  See also getting started guide
extended crashinfo file 46-5
extended-range VLANs
  configuration guidelines 17-8
  configuring 17-8
  creating with an internal VLAN ID 17-18
extended system ID
  MSTP 21-14
  STP 20-3, 20-11
extended universal identifier
  See EUI
Extensible Authentication Protocol over LAN 13-1

F

fa0 interface 1-5
fallback bridging

connecting interfaces with 15-6
VLAN-bridge STP 20-10
Fast Convergence 24-3
FCS bit error rate alarm
configuring 3-8
defined 3-3
FCS error hysteresis threshold 3-2
features, incompatible 29-7
fiber-optic, detecting unidirectional links 33-1
files
basic crashinfo
description 46-5
location 46-5
copying A-6
crashinfo, description 46-5
deleting A-6
displaying the contents of A-9
extended crashinfo
description 46-5
location 46-5
tar
creating A-7
displaying the contents of A-7
extracting A-8
image file format A-23
file system
displaying available file systems A-1
displaying file information A-4
local file system names A-1
network file system names A-6
setting the default A-3
filtering
non-IP traffic 37-11
show and more command output 2-10
filtering show and more command output 2-10
filters, IP
See ACLs, IP
flash device, number of A-1
flexible authentication ordering
overview 13-27
Flex Link Multicast Fast Convergence 24-3
Flex Links
configuration guidelines 24-6
configuring preferred VLAN 24-12
default configuration 24-5
description 24-1
link load balancing 24-2
VLANs 24-2
flow-based packet classification 1-10
flowcharts
QoS classification 38-12
QoS egress queueing and scheduling 38-23
QoS ingress queueing and scheduling 38-21
QoS policing and marking 38-16
flowcontrol
described 15-9
FTP
configuration files
downloading A-14
overview A-13
preparing the server A-13
uploading A-15
image files
deleting old image A-30
downloading A-29
preparing the server A-28
uploading A-30

G

general query 24-10
Generating IGMP Reports 24-3
get-bulk-request operation 36-4
get-next-request operation 36-4, 36-5
get-request operation 36-4, 36-5
get-response operation 36-4
global configuration mode 2-2
global leave, IGMP 28-7

global status monitoring alarms 3-2
guest VLAN and 802.1x 13-20
GUIs
See device manager and Network Assistant

H

hardware limitations and Layer 3 interfaces 15-12
help, for the command line 2-3
hierarchical policy maps 38-14
configuration guidelines 38-5
configuring 38-29, 38-42
described 38-17
history
changing the buffer size 2-6
described 2-6
disabling 2-7
recalling commands 2-6
history table, level and number of syslog messages 35-9
host names, in clusters 6-11
hosts, limit on dynamic ports 17-16
HP OpenView 1-3
HTTP over SSL
see HTTPS
HTTPS 12-22
configuring 12-42
self-signed certificate 12-23
HTTP secure server 12-22

I

ICMP
IPv6 42-3
time-exceeded messages 46-4
traceroute and 46-4
unreachable messages 37-9
unreachables and ACLs 37-10
ICMP Echo operation
IP SLAs 45-6
ICMP ping
executing 46-9
overview 46-2
ICMPv6 42-3
IEEE 802.1D
See STP
IEEE 802.1p 19-1
IEEE 802.1Q
and trunk ports 15-4
configuration limitations 17-10
native VLAN for untagged traffic 17-12, 17-20
IEEE 802.1s
See MSTP
IEEE 802.1w
See RSTP
IEEE 802.3ad
See EtherChannel
IEEE 802.3x flow control 15-9, 15-16
ifIndex values, SNMP 36-6
IFS 1-5
IGMP
configurable leave timer
described 28-5
flooded multicast traffic
controlling the length of time 28-7
disabling on an interface 28-7
global leave 28-7
query solicitation 28-7
recovering from flood mode 28-7
joining multicast group 28-3
join messages 28-3
leave processing, enabling 44-6, 44-8
leaving multicast group 28-5
queries 28-4
report suppression
described 28-6
supported versions 28-2
support for 1-2

IGMP filtering
configuring 28-13
default configuration 28-12
described 28-12
support for 1-2
IGMP groups
configuring filtering 28-13
setting the maximum number 28-1
IGMP helper 1-2
IGMP Immediate Leave
described 28-5
IGMP profile
applying 28-13
configuration mode 28-13
IGMP snooping
and address aliasing 28-2
default configuration 28-6, 44-5
definition 28-2
enabling and disabling 28-14, 44-6
Immediate Leave 28-5
method 28-6
monitoring 44-9
querier
configuring 28-7, 28-16
supported versions 28-2
support for 1-2
IGMP throttling
configuring 28-13
default configuration 28-12
described 28-12
Immediate Leave, IGMP 28-5
enabling 44-6, 44-8
inaccessible authentication bypass 13-22
support for multiauth ports 13-22
initial configuration
defaults 1-11
Express Setup 1-2
interface
range macros 15-14
interface command 15-6 to 15-13
interface configuration mode 2-3
interfaces
auto-MDIX, configuring 15-10, 15-17
configuration guidelines
duplex and speed 15-9
configuring
procedure 15-13
counters, clearing 15-19
default configuration 15-8
descriptive name, adding 15-17
displaying information about 15-18
flow control 15-9, 15-16
management 1-3
physical, identifying 15-6
range of 15-13
restarting 15-19
shutting down 15-19
speed and duplex, configuring 15-16
supported 15-6
types of 15-1
interfaces range macro command 15-14
interface types 15-6
Internet Protocol version 6
See IPv6
inter-VLAN routing 41-2
Intrusion Detection System
See IDS appliances
inventory management TLV 31-3, 31-5
IP ACLs
for QoS classification 38-13
implicit deny 37-7, 37-11
implicit masks 37-11
named 37-7
undefined 37-10
IP addresses
128-bit 42-2
candidate or member 6-2, 6-11
classes of 41-3

cluster access 6-3
command switch 6-1, 6-11
discovering 7-8
for IP routing 41-3
IPv6 42-2
standby command switch 6-11
See also IP information
ip igmp profile command 28-13
IP information
assigned
manually 4-15
through DHCP-based autoconfiguration 4-4
IP multicast routing
and IGMP snooping 28-2
IP phones
and QoS 19-1
automatic classification and queueing 39-2
ensuring port security with QoS 38-26
trusted boundary for QoS 38-26
IP Port Security for Static Hosts
on a PVLAN host port 27-5
IP precedence 38-2
IP-precedence-to-DSCP map for QoS 38-48
IP protocols
in ACLs 37-13
IP routing
connecting interfaces with 15-6
enabling 41-3
IP Service Level Agreements
See IP SLAs
IP service levels, analyzing 45-1
IP SLAs
benefits 45-2
Control Protocol 45-3
definition 45-1
ICMP echo operation 45-6
measuring network performance 45-3
operation 45-3
responder
described 45-3
response time 45-4
scheduling 45-4
SNMP support 45-2
supported metrics 45-2
threshold monitoring 45-5
UDP jitter operation 45-5
IP source guard
and 802.1x 27-4
and DHCP snooping 27-2
and EtherChannels 27-3
and port security 27-3
and private VLANs 27-3
and routed ports 27-3
and TCAM entries 27-4
and trunk interfaces 27-3
and VRF 27-4
binding configuration
automatic 27-2
manual 27-2
binding table 27-2
configuration guidelines 27-3
described 27-1
filtering
source IP address 27-2
source IP and MAC address 27-2
source IP address filtering 27-2
source IP and MAC address filtering 27-2
IP traceroute
executing 46-10
overview 46-4
IP unicast routing
assigning IP addresses to Layer 3 interfaces 41-4
configuring static routes 41-4
enabling 41-3
inter-VLAN 41-2
IP addressing
classes 41-3
configuring 41-3

steps to configure 41-3
subnet mask 41-3
with SVIs 41-3
IPv4 ACLs
applying to interfaces 37-9, 37-17
extended, creating 37-6, 37-13
named 37-7, 37-15
standard, creating 37-11
IPv4 and IPv6
dual protocol stacks 42-4
IPv6
addresses 42-2
address formats 42-2
applications 42-4
assigning address 42-7
autoconfiguration 42-4
default router preference (DRP) 42-4
defined 42-1
forwarding 42-7
ICMP 42-3
neighbor discovery 42-3
SDM templates 11-3, 44-1
Stateless Autoconfiguration 42-4
supported features 42-2
understanding static routes 42-5

J

join messages, IGMP 28-3

K

KDC
described 12-17
See also Kerberos
Kerberos
authenticating to
boundary switch 12-19
KDC 12-19
network services 12-20
configuring 12-20
credentials 12-17
cryptographic software image 12-17
described 12-17
KDC 12-17
operation 12-19
realm 12-18
server 12-19
support for 1-9
terms 12-18
TGT 12-19
tickets 12-17
key distribution center
See KDC

L

LACP
See EtherChannel
Layer 2 frames, classification with CoS 38-2
Layer 2 interfaces, default configuration 15-15
Layer 2 traceroute
and ARP 46-3
and CDP 46-3
broadcast traffic 46-3
described 46-3
IP addresses and subnets 46-3
MAC addresses and VLANs 46-3
multicast traffic 46-3
multiple devices on a port 46-4
unicast traffic 46-3
usage guidelines 46-3
Layer 3 interfaces
assigning IP addresses to 41-4
changing from Layer 2 mode 41-4
Layer 3 packets, classification methods 38-2
LDAP 5-3

Leaking IGMP Reports 24-4
LEDs, switch
See hardware installation guide
lightweight directory access protocol
See LDAP
line configuration mode 2-3
Link Aggregation Control Protocol
See EtherChannel
link failure, detecting unidirectional 21-7
link fault alarm 3-3
link integrity, verifying with REP 23-4
link local unicast addresses 42-3
link redundancy
See Flex Links
links, unidirectional 33-1
link-state tracking
configuring 43-4
described 43-1
LLDP
configuring
characteristics 31-5
default configuration 31-4
enabling 31-5
monitoring and maintaining 31-8, 31-9
supported TLVs 31-2
switch stack considerations 31-2
transmission timer and holdtime, setting 31-5
LLDP-MED
configuring
TLVs 31-6
monitoring and maintaining 31-8, 31-9
overview 31-2
supported TLVs 31-2
local SPAN 30-2
location TLV 31-3, 31-5
logging messages, ACL 37-6
login authentication
with RADIUS 12-15, 12-36
with TACACS+ 12-7, 12-31
login banners 7-4
loop guard
described 22-8
support for 1-6
LRE profiles, considerations in switch clusters 6-13

M

MAB
See MAC authentication bypass
MAB aging timer 1-7
MAB inactivity timer
default setting 13-31
range 13-34
MAC/PHY configuration status TLV 31-2
MAC addresses
aging time 7-6
and VLAN association 7-5
building the address table 7-5
discovering 7-8
dynamic
learning 7-5
in ACLs 37-11
static
allowing 7-8
characteristics of 7-6
MAC address learning 1-4
MAC address notification, support for 1-11
MAC address-table move update
configuration guidelines 24-6
configuring 24-8
default configuration 24-5
description 24-4
MAC address-to-VLAN mapping 17-14
MAC authentication bypass 13-33
overview 13-14
See MAB
MAC extended access lists
applying to Layer 2 interfaces 37-11

creating 37-11
defined 37-11
for QoS classification 38-10
magic packet 13-24
manageability features 1-4
management access
in-band
browser session 1-5
CLI session 1-5
device manager 1-5
SNMP 1-5
out-of-band console port connection 1-5
management address TLV 31-2
management options
CLI 2-1
CNS 5-2
overview 1-3
management VLAN
considerations in switch clusters 6-8
discovery through different management VLANs 6-8
manual preemption, REP, configuring 23-12
mapping tables for QoS
configuring
CoS-to-DSCP 38-47
DSCP 38-47
DSCP-to-CoS 38-48
DSCP-to-DSCP-mutation 38-30, 38-49
IP-precedence-to-DSCP 38-48
policed-DSCP 38-48
described 38-18
marking
action with aggregate policers 38-46
described 38-4, 38-14
matching, IPv4 ACLs 37-5
maximum number of allowed devices, port-based authentication 13-34
MDA
configuration guidelines 13-10 to 13-11
described 1-8, 13-10
exceptions with authentication process 13-4
membership mode, VLAN port 17-3
member switch
automatic discovery 6-5
defined 6-3
managing 6-13
passwords 6-11
recovering from lost connectivity 46-9
requirements 6-2
See also candidate switch, cluster standby group, and standby command switch
messages, to users through banners 7-4
MIBs
overview 36-2
SNMP interaction with 36-5
mirroring traffic for analysis 30-2
mismatches, autonegotiation 46-1
monitoring
alarms 3-9
cables for unidirectional links 33-1
CDP 32-3
features 1-11
IGMP
snooping 44-9
multicast router interfaces 28-20, 44-9
network traffic for analysis with probe 30-2
port
blocking 29-16
protection 29-16
PTP 8-3, 8-4
SFP status 46-13
speed and duplex mode 15-16
traffic flowing among switches 34-1
VTP 18-14
mrouter Port 24-3
mrouter port 24-10
MSTP
boundary ports
configuration guidelines 21-14

described 21-5
BPDU filtering
described 22-3
BPDU guard
described 22-2
CIST, described 21-3
CIST regional root 21-3
CIST root 21-4
configuration guidelines 21-13
configuring
link type for rapid convergence 21-15
MST region 21-16
neighbor type 21-15
path cost 21-15
port priority 21-15
root switch 21-14, 21-17
secondary root switch 21-15
CST
defined 21-3
operations between regions 21-3
default configuration 21-13
enabling the mode 21-16
EtherChannel guard
described 22-7
extended system ID
effects on root switch 21-14
effects on secondary root switch 21-15
unexpected behavior 21-14
IEEE 802.1s
implementation 21-6
port role naming change 21-6
terminology 21-4
instances supported 20-9
interface state, blocking to forwarding 22-1
interoperability and compatibility among modes 20-10
interoperability with IEEE 802.1D
described 21-8
restarting migration process 21-16
IST
defined 21-2
master 21-3
operations within a region 21-3
loop guard
described 22-8
mapping VLANs to MST instance 21-16
MST region
CIST 21-3
configuring 21-16
described 21-2
hop-count mechanism 21-5
IST 21-2
supported spanning-tree instances 21-2
optional features supported 1-6
overview 21-2
Port Fast
described 22-1
preventing root switch selection 22-7
root guard
described 22-7
root switch
configuring 21-14
effects of extended system ID 21-14
unexpected behavior 21-14
shutdown Port Fast-enabled port 22-2
multiauth
support for inaccessible authentication bypass 13-22
multiauth mode
See multiple-authentication mode
multicast groups
Immediate Leave 28-5
joining 28-3
leaving 28-5
static joins 44-7
multicast router interfaces, monitoring 28-20, 44-9
multicast router ports, adding 28-14, 44-6
multicast storm 29-1
multicast storm-control command 29-9

multicast television application 28-9
multicast VLAN 28-8
Multicast VLAN Registration
See MVR
multidomain authentication
See MDA
multiple authentication 13-11
multiple authentication mode
configuring 13-38
MVR
and address aliasing 28-11
and IGMPv3 28-11
default configuration 28-11
described 28-8
example application 28-9
modes 28-17
multicast television application 28-9
setting global parameters 28-16

N

NAC
AAA down policy 1-9
critical authentication 13-22, 13-44
IEEE 802.1x authentication using a RADIUS server 13-46
IEEE 802.1x validation using RADIUS server 13-46
inaccessible authentication bypass 1-9, 13-44
Layer 2 IEEE 802.1x validation 1-8, 13-27, 13-46
Layer 2 IP validation 1-9
named IPv4 ACLs 37-7
NameSpace Mapper
See NSM
native VLAN
configuring 17-12, 17-20
default 17-12
NEAT
configuring 13-47
overview 13-28
neighbor discovery, IPv6 42-3
neighbor offset numbers, REP 23-5
Network Admission Control
NAC
Network Assistant
benefits 1-2
described 1-3
upgrading a switch A-23
network configuration examples
increasing network performance 1-14
providing network services 1-14
network design
performance 1-14
services 1-14
Network Edge Access Topology
See NEAT
network management
CDP 32-1
RMON 34-1
SNMP 36-1
network performance, measuring with IP SLAs 45-3
network policy TLV 31-2, 31-5
Network Time Protocol
See NTP
no commands 2-4
nonhierarchical policy maps
configuration guidelines 38-5
described 38-15
non-IP traffic filtering 37-11
nontrunking mode 17-10
normal-range VLANs 17-4
configuration guidelines 17-6
configuring 17-4
no switchport command 15-3
NSM 5-4
NTP
associations
defined 7-2
overview 7-2

stratum 7-2
support for 1-4
time
services 7-2
synchronizing 7-2

O

off mode, VTP 18-3
open1x
configuring 13-50
open1x authentication
overview 13-28
Open DeviceNet Vendors Association (ODVA) 10-1
optimizing system resources 11-1
options, management 1-3
out-of-profile markdown 1-10

P

packet modification, with QoS 38-25
PAgP
See EtherChannel
passwords
default configuration 12-2
disabling recovery of 12-3, 12-27
encrypting 12-3, 12-27
for security 1-7
in clusters 6-12
overview 12-2
recovery of 46-8
setting
enable 12-26
enable secret 12-3, 12-27
Telnet 12-28
with usernames 12-4
VTP domain 18-10
path cost
MSTP 21-15
STP 20-13, 20-16
performance, network design 1-14
performance features 1-2
persistent self-signed certificate 12-23
per-user ACLs and Filter-Ids 13-7
per-VLAN spanning-tree plus
See PVST+
physical ports 15-2
PIM-DVMRP, as snooping method 28-7
ping
character output description 46-10
executing 46-9
overview 46-2
policed-DSCP map for QoS 38-48
policers
configuring
for each matched traffic class 38-28
for more than one traffic class 38-46
described 38-4
displaying 38-57
number of 38-6
types of 38-15
policing
described 38-4
hierarchical
See hierarchical policy maps
token-bucket algorithm 38-15
policy maps for QoS
characteristics of 38-28
described 38-13
displaying 38-57
hierarchical 38-14
hierarchical on SVIs
configuration guidelines 38-5
configuring 38-29, 38-42
described 38-17
nonhierarchical on physical ports
configuration guidelines 38-5

described 38-15
port ACLs
defined 37-2
types of 37-2
Port Aggregation Protocol
See EtherChannel
port-based authentication
accounting 13-13
authentication server
defined 13-2, 14-2
RADIUS server 13-2
client, defined 13-2, 14-2
configuring
guest VLAN 13-42
host mode 13-38
inaccessible authentication bypass 13-44
RADIUS server parameters on the switch 13-4, 13-36, 14-10
restricted VLAN 13-43
default configuration 13-30, 14-10
described 13-1
device roles 13-2, 14-2
downloadable ACLs and redirect URLs
overview 13-18 to 13-19
EAPOL-start frame 13-4
EAP-request/identity frame 13-4
EAP-response/identity frame 13-4
enabling
802.1X authentication 14-10, 14-12
encapsulation 13-2
flexible authentication ordering
overview 13-27
guest VLAN
configuration guidelines 13-21, 13-22
described 13-20
host mode 13-9
inaccessible authentication bypass
configuring 13-44
described 13-22
guidelines 13-33
initiation and message exchange 13-4
magic packet 13-24
maximum number of allowed devices per port 13-34
method lists 13-34
multiple authentication 13-11
per-user ACLs
AAA authorization 13-34
configuration tasks 13-17
described 13-17
RADIUS server attributes 13-17
ports
authorization state and dot1x port-control command 13-9
authorized and unauthorized 13-9
voice VLAN 13-23
port security
described 13-24
readiness check
configuring 13-14
described 13-14
resetting to default values 13-51
switch
as proxy 13-2, 14-3
RADIUS client 13-2
switch supplicant
configuring 13-47
overview 13-28
user distribution
guidelines 13-26
overview 13-26
VLAN assignment
AAA authorization 13-34
characteristics 13-15
configuration tasks 13-16
described 13-15
voice aware 802.1x security
configuring 13-16
described 13-16

voice VLAN
described 13-23
PVID 13-23
VVID 13-23
wake-on-LAN, described 13-24
with ACLs and RADIUS Filter-Id attribute 13-29
port-based authentication configuration process 13-34
port-based authentication methods, supported 13-6
port blocking 1-2, 29-4, 29-11
port-channel
See EtherChannel
port description TLV 31-2
Port Fast
described 22-1
mode, spanning tree 17-15
support for 1-6
port membership modes, VLAN 17-3
port not forwarding alarm 3-3
port not operating alarm 3-3
port priority
MSTP 21-15
STP 20-12, 20-16
ports
access 15-3
blocking 29-4, 29-11
dual-purpose uplink 15-4
dynamic access 17-4
protected 29-3
REP 23-6
routed 15-3
secure 29-4
static-access 17-3, 17-7, 17-17
switch 15-2
trunks 17-3, 17-9
VLAN assignments 17-7, 17-17
port security
aging 29-8, 29-15
and QoS trusted boundary 38-26, 38-34
described 29-4
displaying 29-16
on trunk ports 29-13
sticky learning 29-5
violations 29-5
port-shutdown response, VMPS 17-14
port status monitoring alarms
FCS bit error rate alarm 3-3
link fault alarm 3-3
port not forwarding alarm 3-3
port not operating alarm 3-3
port VLAN ID TLV 31-2
power management TLV 31-2, 31-5
preempt delay time, REP 23-5
preemption, default configuration 24-5
preemption delay, default configuration 24-5
preferential treatment of traffic
See QoS
preventing unauthorized access 12-2
primary edge port, REP 23-4
primary links 24-2
priority
overriding CoS 19-4
trusting CoS 19-4
private VLAN edge ports
See protected ports
privileged EXEC mode 2-2
privilege levels
changing the default for lines 12-29
command switch 6-13
exiting 12-30
logging into 12-30
mapping on member switches 6-13
overview 12-2, 12-4
setting a command with 12-29
PROFINET
configuring 9-4
default configuration 9-4
protected ports 1-7, 29-3
protocol storm protection 29-8

proxy reports 24-3
pruning, VTP
enabling
in VTP domain 18-13
on a port 17-19
examples 18-8
overview 18-7
pruning-eligible list
changing 17-19
for VTP pruning 18-7
PTP
configuring 8-3
default configuration 8-2
displaying configuration 8-3, 8-4
PVST+
described 20-9
IEEE 802.1Q trunking interoperability 20-10
instances supported 20-9

Q

QoS
and MQC commands 38-2, 39-2
auto-QoS
categorizing traffic 39-3
disabling 39-8
displaying generated commands 39-8
effects on running configuration 39-7
egress queue defaults 39-3
enabling for VoIP 39-8
ingress queue defaults 39-3
list of generated commands 39-4
basic model 38-4
classification
class maps, described 38-13
defined 38-4
DSCP transparency, described 38-27, 38-34
flowchart 38-12
forwarding treatment 38-3
in frames and packets 38-3
IP ACLs, described 38-11, 38-13
MAC ACLs, described 38-10, 38-13
options for IP traffic 38-10
options for non-IP traffic 38-10
policy maps, described 38-13
trust DSCP, described 38-10
trusted CoS, described 38-10
trust IP precedence, described 38-10
class maps
configuring 38-38
displaying 38-56
configuring
aggregate policers 38-46
default port CoS value 38-33
DSCP maps 38-47
DSCP transparency 38-27, 38-34
DSCP trust states bordering another domain 38-27, 38-35
egress queue characteristics 38-31, 38-52
ingress queue characteristics 38-30, 38-49
IP standard ACLs 38-36
policy maps, hierarchical 38-29, 38-42
port trust states within the domain 38-26, 38-33
trusted boundary 38-26, 38-34
default auto configuration 39-3
default standard configuration 38-6
displaying statistics 38-57
DSCP transparency 38-27, 38-34
egress queues
allocating buffer space 38-31, 38-52
buffer allocation scheme, described 38-23
configuring shaped weights for SRR 38-54
configuring shared weights for SRR 38-55
described 38-4
flowchart 38-23
mapping DSCP or CoS values 38-53
scheduling, described 38-4
setting WTD thresholds 38-31, 38-52

WTD, described 38-24
enabling globally 38-32
flowcharts
classification 38-12
egress queueing and scheduling 38-23
ingress queueing and scheduling 38-21
policing and marking 38-16
implicit deny 38-13
ingress queues
allocating bandwidth 38-51
allocating buffer space 38-50
buffer and bandwidth allocation, described 38-22
configuring shared weights for SRR 38-51
configuring the priority queue 38-30, 38-51
described 38-4
displaying the threshold map 38-57
flowchart 38-21
mapping DSCP or CoS values 38-49
priority queue, described 38-22
scheduling, described 38-4
setting WTD thresholds 38-49
WTD, described 38-22
IP phones
automatic classification and queueing 39-2
detection and trusted settings 38-26, 39-2
limiting bandwidth on egress interface 38-56
mapping tables
CoS-to-DSCP 38-47
displaying 38-57
DSCP-to-CoS 38-48
DSCP-to-DSCP-mutation 38-30, 38-49
IP-precedence-to-DSCP 38-48
policed-DSCP 38-48
types of 38-18
marked-down actions 38-41, 38-44
marking, described 38-4, 38-14
packet modification 38-25
policers
configuring 38-41, 38-44
described 38-14
displaying 38-57
number of 38-6
types of 38-15
policies, attaching to an interface 38-14
policing
described 38-4, 38-14
token bucket algorithm 38-15
policy maps
characteristics of 38-28
displaying 38-57
hierarchical 38-14
hierarchical on SVIs 38-29
nonhierarchical on physical ports 38-28
QoS label, defined 38-4
queues
configuring egress characteristics 38-31, 38-52
configuring ingress characteristics 38-30, 38-49
high priority (expedite) 38-25, 38-56
location of 38-19
SRR, described 38-20
WTD, described 38-19
rewrites 38-25
support for 1-10
trust states
bordering another domain 38-27, 38-35
described 38-10
trusted device 38-26, 38-34
within the domain 38-26, 38-33
quality of service
See QoS
queries, IGMP 28-4
query solicitation, IGMP 28-7

R

RADIUS
attributes
vendor-proprietary 12-17, 12-38

vendor-specific 12-16
configuring
accounting 12-16, 12-37
authentication 12-15, 12-36
authorization 12-16, 12-37
communication, global 12-15, 12-37
communication, per-server 12-14, 12-15
multiple UDP ports 12-14
default configuration 12-10
defining AAA server groups 12-15, 12-35
identifying the server 12-14
in clusters 6-12
limiting the services to the user 12-16
method list, defined 12-15
operation of 12-9
overview 12-8
suggested network environments 12-8
support for 1-9
tracking services accessed by user 12-16, 12-37
RADIUS Change of Authorization 12-10
range
macro 15-14
rapid convergence 21-9
rapid per-VLAN spanning-tree plus
See rapid PVST+
rapid PVST+
described 20-9
IEEE 802.1Q trunking interoperability 20-10
instances supported 20-9
Rapid Spanning Tree Protocol
See RSTP
rcommand command 6-13
RCP
configuration files
downloading A-17
overview A-16
preparing the server A-16
uploading A-18
image files
deleting old image A-34
downloading A-33
preparing the server A-32
uploading A-34
readiness check
port-based authentication
configuring 13-14
described 13-14
reconfirmation interval, VMPS, changing 17-16
redirect URL 13-18, 13-19, 13-48
redundancy
EtherChannel 40-2
STP
backbone 20-7
path cost 17-13
port priority 17-13
reloading software 4-11
Remote Authentication Dial-In User Service
See RADIUS
Remote Copy Protocol
See RCP
Remote SPAN
See RSPAN
remote SPAN 30-2
REP
administrative VLAN 23-9
age timer 23-8
and STP 23-6
configuration guidelines 23-7
convergence 23-4
default configuration 23-7
manual preemption, configuring 23-12
neighbor offset numbers 23-5
open segment 23-2
ports 23-6
preempt delay time 23-5
primary edge port 23-4
ring segment 23-2
secondary edge port 23-4

segments 23-2
characteristics 23-3
SNMP traps, configuring 23-12
supported interfaces 23-2
triggering VLAN load balancing 23-6
verifying link integrity 23-4
VLAN blocking 23-14
VLAN load balancing 23-4
report suppression, IGMP
described 28-6
resequencing ACL entries 37-7
resetting a UDLD-shutdown interface 33-5
Resilient Ethernet Protocol
See REP
responder, IP SLAs
described 45-3
response time, measuring with IP SLAs 45-4
restricted VLAN
configuring 13-43
described 13-21
using with IEEE 802.1x 13-21
restricting access
overview 12-2
passwords and privilege levels 12-2
RADIUS 12-8
TACACS+ 12-5
RFC
1112, IP multicast and IGMP 28-2
1157, SNMPv1 36-2
1166, IP addresses 41-3
1305, NTP 7-2
1757, RMON 34-2
1901, SNMPv2C 36-2
1902 to 1907, SNMPv2 36-2
2236, IP multicast and IGMP 28-2
2273-2275, SNMPv3 36-2
RFC 5176 Compliance 12-10
RMON
enabling alarms and events 34-3
groups supported 34-2
overview 34-1
statistics
collecting group Ethernet 34-4
collecting group history 34-4
support for 1-11
root guard
described 22-7
support for 1-6
root switch
MSTP 21-14, 21-17
STP 20-11, 20-15
routed ports
defined 15-3
in switch clusters 6-9
IP addresses on 15-12
RSPAN
characteristics 30-7
configuration guidelines 30-9
default configuration 30-10
defined 30-2
destination ports 30-6
monitored ports 30-5
monitoring ports 30-6
overview 1-11, 30-1
received traffic 30-4
sessions
defined 30-3
source ports 30-5
transmitted traffic 30-4
VLAN-based 30-6
RSTP
active topology 21-8
BPDU
format 21-11
processing 21-12
designated port, defined 21-8
designated switch, defined 21-8
interoperability with IEEE 802.1D

described 21-8
restarting migration process 21-16
topology changes 21-12
overview 21-8
port roles
described 21-8
synchronized 21-10
proposal-agreement handshake process 21-9
rapid convergence
described 21-9
edge ports and Port Fast 21-9
point-to-point links 21-9, 21-15
root ports 21-9
root port, defined 21-8
See also MSTP
running configuration
replacing A-19, A-20
rolling back A-19, A-20
running configuration, saving 4-17

S

scheduled reloads 4-11
scheduling, IP SLAs operations 45-4
SCP
and SSH 12-25
SD flash memory card A-2
SDM
templates
number of 11-2
SDM template
configuring 11-4
dual IPv4 and IPv6 11-3
types of 11-2
secondary edge port, REP 23-4
Secure Copy Protocol
Secure Digital flash memory card
See SD flash memory card
secure HTTP client
configuring 12-44
secure MAC addresses
maximum number of 29-5
types of 29-4
secure ports, configuring 29-4
secure remote connections 12-21
Secure Shell
See SSH
Secure Socket Layer
See SSL
security, port 29-4
security features 1-7
See SCP
sequence numbers in log messages 35-8
server mode, VTP 18-3
service-provider network, MSTP and RSTP 21-1
set-request operation 36-5
setting the FCS error hysteresis threshold 3-8
severity levels, defining in system messages 35-8
SFPs
monitoring status of 46-13
security and identification 46-2
status, displaying 46-13
shaped round robin
See SRR
show access-lists hw-summary command 37-10
show and more command output, filtering 2-10
show cdp traffic command 32-3
show cluster members command 6-13
show forward command 46-14
show interfaces command 15-16
show interfaces switchport 24-9
show lldp traffic command 31-9
show platform forward command 46-14
shutdown command on interfaces 15-19
Simple Network Management Protocol
See SNMP
small-frame arrival rate, configuring 29-10
Smartports macros

applying global parameter values 16-3
configuration guidelines 16-2
default configuration 16-1
tracing 16-2
SNAP 32-1
SNMP
accessing MIB variables with 36-5
agent
described 36-4
disabling 36-8
and IP SLAs 45-2
authentication level 36-11
community strings
configuring 36-6, 36-9
for cluster switches 36-4
overview 36-4
configuration examples 36-15
default configuration 36-8
engine ID 36-1
groups 36-1, 36-10
host 36-1
ifIndex values 36-6
in-band management 1-5
in clusters 6-12
informs
and trap keyword 36-6, 36-12
described 36-5
differences from traps 36-5
limiting access by TFTP servers 36-15
limiting system log messages to NMS 35-9
manager functions 1-3, 36-4
managing clusters with 6-14
notifications 36-5
overview 36-2, 36-5
security levels 36-3
setting CPU threshold notification 36-14
system contact and location 36-14
trap manager, configuring 36-12
traps
described 36-4, 36-5
differences from informs 36-5
enabling 36-6, 36-12
enabling MAC address notification 7-14, 7-15
overview 36-2, 36-5
types of 36-7
users 36-1, 36-10
versions supported 36-2
SNMP and Syslog Over IPv6 42-5
SNMP traps
REP 23-12
SNMPv1 36-3
SNMPv2C 36-3
SNMPv3 36-3
snooping, IGMP 28-2
software images
location in flash A-23
recovery procedures 46-7
tar file format, described A-23
See also downloading and uploading
source addresses
in IPv4 ACLs 37-13
source-and-destination-IP address based forwarding, EtherChannel 40-9
source-and-destination MAC address forwarding, EtherChannel 40-8
source-IP address based forwarding, EtherChannel 40-8
source-MAC address forwarding, EtherChannel 40-8
SPAN
configuration guidelines 30-9
default configuration 30-10
destination ports 30-6
monitored ports 30-5
monitoring ports 30-6
overview 1-11, 30-1
ports, restrictions 29-7
received traffic 30-4
sessions
defined 30-3

source ports 30-5
transmitted traffic 30-4
VLAN-based 30-6
spanning tree and native VLANs 17-10
Spanning Tree Protocol
See STP
SRR
configuring
shaped weights on egress queues 38-54
shared weights on egress queues 38-55
shared weights on ingress queues 38-51
described 38-20
shaped mode 38-20
shared mode 38-20
support for 1-10, 1-11
SSH
cryptographic software image 12-1, 12-21
described 1-5, 12-21
encryption methods 12-21
user authentication methods, supported 12-21
SSL
configuring a secure HTTP client 12-44
configuring a secure HTTP server 12-42
cryptographic software image 12-22
described 12-22
standby command switch
configuring
defined 6-3
requirements 6-2
See also cluster standby group and HSRP
standby group, cluster
See cluster standby group and HSRP
standby links 24-2
startup configuration
booting
specific image 4-17
clearing A-19
configuration file
specifying the filename 4-15
static access ports
assigning to VLAN 17-7, 17-17
defined 15-3, 17-3
static addresses
See addresses
static MAC addressing 1-7
static routes
configuring 41-4
understanding 42-5
static VLAN membership 17-2
statistics
interface 15-18
QoS ingress and egress 38-57
RMON group Ethernet 34-4
RMON group history 34-4
VTP 18-14
sticky learning 29-5
storm control
configuring 29-3, 29-9
displaying 29-16
support for 1-2
thresholds 29-1
STP
accelerating root port selection 22-4
and REP 23-6
BackboneFast
described 22-5
BPDU filtering
described 22-3
BPDU guard
described 22-2
BPDU message exchange 20-2
configuration guidelines 20-13
configuring
path cost 20-13, 20-16
port priority 20-12, 20-16
root switch 20-11, 20-15
secondary root switch 20-12, 20-16
spanning-tree mode 20-14

switch priority 20-17
default configuration 20-11
default optional feature configuration 22-9
designated port, defined 20-3
designated switch, defined 20-3
detecting indirect link failures 22-5
disabling 20-11
EtherChannel guard
described 22-7
extended system ID
effects on root switch 20-11
effects on the secondary root switch 20-12
overview 20-3
unexpected behavior 20-12
features supported 1-6
IEEE 802.1D and bridge ID 20-3
IEEE 802.1D and multicast addresses 20-8
IEEE 802.1t and VLAN identifier 20-4
inferior BPDU 20-3
instances supported 20-9
interface state, blocking to forwarding 22-1
interface states
blocking 20-5
disabled 20-6
forwarding 20-5, 20-6
learning 20-6
listening 20-6
overview 20-4
interoperability and compatibility among modes 20-10
limitations with IEEE 802.1Q trunks 20-10
load sharing
overview 17-12
using path costs 17-13
using port priorities 17-13
loop guard
described 22-8
modes supported 20-9
multicast addresses, effect of 20-8
optional features supported 1-6
path costs 17-13
Port Fast
described 22-1
preventing root switch selection 22-7
protocols supported 20-9
redundant connectivity 20-7
root guard
described 22-7
root port, defined 20-3
root switch
configuring 20-12
effects of extended system ID 20-3, 20-11
election 20-3
unexpected behavior 20-12
shutdown Port Fast-enabled port 22-2
superior BPDU 20-3
UplinkFast
described 22-3
VLAN-bridge 20-10
stratum, NTP 7-2
subnet mask 41-3
success response, VMPS 17-14
summer time 7-10
SunNet Manager 1-3
supported port-based authentication methods 13-6
SVI autostate exclude
configuring 15-10
SVIs
and IP unicast routing 41-3
connecting VLANs 15-5
routing between VLANs 17-2
switch 42-2
switch boot process 4-1
switch console port 1-5
Switch Database Management
See SDM
Switched Port Analyzer
See SPAN

switched ports 15-2
switch information assignment 4-4
switchport backup interface 24-4, 24-10
switchport block multicast command 29-11
switchport block unicast command 29-11
switchport command 15-8
switchport protected command 29-11
switch priority
  STP 20-17
switch software features 1-1
system capabilities TLV 31-2
system clock
  configuring
    daylight saving time 7-10
    manually 7-9
    summer time 7-10
    time zones 7-9
  overview 7-1
  See also NTP
system description TLV 31-2
system message logging
  default configuration 35-5
  defining error message severity levels 35-8
  disabling 35-5
  facility keywords, described 35-4
  level keywords, described 35-3
  limiting messages 35-9
  message format 35-2
  overview 35-1
  sequence numbers, enabling and disabling 35-8
  setting the display destination device 35-6
  synchronizing log messages 35-2, 35-7
  syslog facility 1-11
  time stamps, enabling and disabling 35-8
  UNIX syslog servers
    configuring the daemon 35-4
    configuring the logging facility 35-10
    facilities supported 35-4
system name
  manual configuration 7-11
  See also DNS
system name TLV 31-2
system prompt, default setting 7-5
system resources, optimizing 11-1

T

TACACS+
  accounting, defined 12-6
  authentication, defined 12-6
  authorization, defined 12-6
  configuring
    accounting 12-8, 12-33
    authentication key 12-7, 12-30
    authorization 12-7, 12-33
    login authentication 12-7, 12-31
  default configuration 12-7
  identifying the server 12-7, 12-30
  in clusters 6-12
  limiting the services to the user 12-7
  operation of 12-6
  overview 12-5
  support for 1-9
  tracking services accessed by user 12-8, 12-33
tar files
  creating A-7
  displaying the contents of A-7
  extracting A-8
  image file format A-23
TDR 1-11
Telnet
  accessing management interfaces 2-10
  number of connections 1-5
  setting a password 12-28
temporary self-signed certificate 12-23
Terminal Access Controller Access Control System Plus
  See TACACS+
terminal lines, setting a password 12-28

TFTP
configuration files
downloading A-11
preparing the server A-11
uploading A-12
configuration files in base directory 4-8
configuring for autoconfiguration 4-7
image files
deleting A-26
downloading A-25
preparing the server A-25
uploading A-27
limiting access by servers 36-15
TFTP server 1-4
threshold, traffic level 29-2
threshold monitoring, IP SLAs 45-5
time
See NTP and system clock
Time Domain Reflector
See TDR
time-range command 37-8
time ranges in ACLs 37-8, 37-16
time stamps in log messages 35-8
time zones 7-9
TLVs
defined 31-2
LLDP 31-2
LLDP-MED 31-2
Token Ring VLANs
support for 17-5
VTP support 18-5
ToS 1-10
traceroute, Layer 2
and ARP 46-3
and CDP 46-3
broadcast traffic 46-3
described 46-3
IP addresses and subnets 46-3
MAC addresses and VLANs 46-3
multicast traffic 46-3
multiple devices on a port 46-4
unicast traffic 46-3
usage guidelines 46-3
traceroute command 46-10
See also IP traceroute
traffic
blocking flooded 29-11
fragmented 37-3
unfragmented 37-3
traffic policing 1-10
traffic suppression 29-1
transparent mode, VTP 18-3
trap-door mechanism 4-2
traps
configuring MAC address notification 7-14, 7-15
configuring managers 36-6, 36-12
defined 36-4
enabling 7-14, 7-15, 36-6, 36-12
notification types 36-7
overview 36-2, 36-5
triggering alarm options
configurable relay 3-3
methods 3-3
SNMP traps 3-3
syslog messages 3-4
troubleshooting
CPU utilization 46-6
displaying crash information 46-5
setting packet forwarding 46-14
SFP security and identification 46-2
show forward command 46-14
with CiscoWorks 36-5
with ping 46-2
with system message logging 35-1
with traceroute 46-4
trunk failover
See link-state tracking
trunking encapsulation 1-6

trunk ports
defined 15-4, 17-3
trunks
allowed-VLAN list 17-12, 17-19
load sharing
setting STP path costs 17-13
using STP port priorities 17-13
native VLAN for untagged traffic 17-12, 17-20
parallel 17-13
pruning-eligible list 17-19
to non-DTP device 17-9
trusted boundary for QoS 38-26, 38-34
trusted port states
between QoS domains 38-27, 38-35
classification options 38-10
ensuring port security for IP phones 38-26, 38-34
support for 1-10
within a QoS domain 38-26, 38-33
trustpoints, CA 12-23
twisted-pair Ethernet, detecting unidirectional links 33-1
type of service
See ToS

U

UDLD
default configuration 33-4
echoing detection mechanism 33-3
enabling
globally 33-4
per interface 33-5
link-detection mechanism 33-1
neighbor database 33-2
overview 33-1
resetting an interface 33-5
support for 1-5
UDP jitter operation, IP SLAs 45-5
unauthorized ports with IEEE 802.1x 13-9
unicast MAC address filtering 1-4
and adding static addresses 7-7
and broadcast MAC addresses 7-7
and CPU packets 7-7
and multicast addresses 7-7
and router MAC addresses 7-7
configuration guidelines 7-7
described 7-7
unicast storm 29-1
unicast storm control command 29-9
UNIX syslog servers
daemon configuration 35-4
facilities supported 35-4
message logging configuration 35-10
unrecognized Type-Length-Value (TLV) support 18-5
upgrading software images
See downloading
UplinkFast
described 22-3
uploading
configuration files
preparing A-11, A-13, A-16
reasons for A-9
using FTP A-15
using RCP A-18
using TFTP A-12
image files
preparing A-25, A-28, A-32
reasons for A-23
using FTP A-30
using RCP A-34
using TFTP A-27
user EXEC mode 2-2
username-based authentication 12-4

V

version-dependent transparent mode 18-5
virtual switches and PAgP 40-5
vlan.dat file 17-4

VLAN 1, disabling on a trunk port 17-12
VLAN 1 minimization 17-12
vlan-assignment response, VMPS 17-14
VLAN blocking, REP 23-14
VLAN configuration
at bootup 17-5
saving 17-5
VLAN configuration mode 2-2
VLAN database
and startup configuration file 17-5
VLAN configuration saved in 17-5
VLANs saved in 17-4
VLAN filtering and SPAN 30-6
vlan global configuration command 17-5
VLAN ID, discovering 7-8
VLAN load balancing
REP 23-4
VLAN load balancing, triggering 23-6
VLAN load balancing on flex links 24-2
configuration guidelines 24-6
VLAN management domain 18-2
VLAN membership
modes 17-3
VLAN Query Protocol
See VQP
VLANs
adding 17-17
adding to VLAN database 17-7
aging dynamic addresses 20-8
allowed on trunk 17-12, 17-19
and spanning-tree instances 17-3, 17-6, 17-8
configuration guidelines, extended-range VLANs 17-8
configuration guidelines, normal-range VLANs 17-6
configuring 17-1
configuring IDs 17-8
connecting through SVIs 15-5
default configuration 17-6
deleting 17-7, 17-17
described 15-2, 17-1
extended-range 17-8
features 1-6
illustrated 17-2
internal 17-9
modifying 17-17
multicast 28-8
native, configuring 17-12, 17-20
normal-range 17-4
number supported 1-6
parameters 17-5
port membership modes 17-3
static-access ports 17-7, 17-17
STP and IEEE 802.1Q trunks 20-10
supported 17-2
Token Ring 17-5
traffic between 17-2
VLAN-bridge STP 20-10
VLAN Trunking Protocol
See VTP
VLAN trunks 17-9
VMPS
configuration example 17-24
configuration guidelines 17-15
default configuration 17-15
description 17-14
dynamic port membership
described 17-15
reconfirming 17-16
troubleshooting 17-16
entering server address 17-22
mapping MAC addresses to VLANs 17-14
reconfirmation interval, changing 17-16
voice aware 802.1x security
port-based authentication
configuring 13-16
described 13-16
voice-over-IP 19-1
voice VLAN

Cisco 7960 phone, port connections 19-1
configuration guidelines 19-3
configuring IP phones for data traffic
override CoS of incoming frame 19-4
trust CoS priority of incoming frame 19-4
configuring ports for voice traffic in
802.1p priority tagged frames 19-5
802.1Q frames 19-5
connecting to an IP phone 19-4
default configuration 19-3
described 19-1
displaying 19-6
IP phone data traffic, described 19-3
IP phone voice traffic, described 19-2
VQP 1-6, 17-14
VTP
adding a client to a domain 18-10, 18-13
advertisements 17-11, 18-4
and extended-range VLANs 17-2, 18-2
and normal-range VLANs 17-2, 18-2
configuration
guidelines 18-9
saving 18-9
configuration requirements 18-1, 20-1
configuration revision number
guideline 18-10, 18-13
resetting 18-14
consistency checks 18-5
default configuration 18-9
described 18-2
domain names 18-10
domains 18-2
modes
client 18-3
off 18-3
server 18-3
transparent 18-3
monitoring 18-14
passwords 18-10
pruning
enabling 18-13
examples 18-8
overview 18-7
support for 1-6
pruning-eligible list, changing 17-19
statistics 18-14
support for 1-6
Token Ring support 18-5
transparent mode, configuring 18-4
using 18-2
Version
enabling 18-12
version, guidelines 18-6
Version 1 18-5
Version 2
configuration guidelines 18-6
overview 18-5
Version 3
overview 18-5

W

web authentication 13-14
described 1-7
web-based authentication
customizeable web pages 14-6
description 14-2
web-based authentication, interactions with other features 14-8
weighted tail drop
See WTD
wired location service
configuring 31-7
displaying 31-8, 31-9
location TLV 31-3
understanding 31-3
WTD
described 38-19

setting thresholds
egress queue-sets 38-31, 38-52
ingress queues 38-49
support for 1-10, 1-11

X

Xmodem protocol 46-7